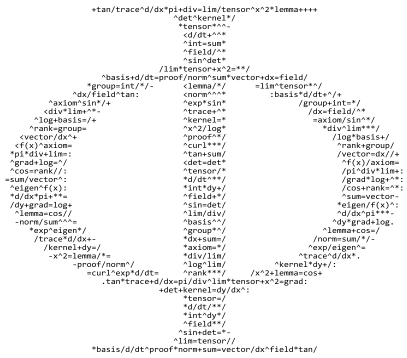
PHI TECHNICAL MONOGRAPHS

Volume IX

Puzzle Geometry and Pinched Disk Models

Planar decompositions, quotients and quantitative models



PHI research collective

Preface and reading contract

The boundary theory of the preceding volume distinguishes a point of a continuum from an approach to that point. The present volume asks how much of a continuum can be recovered from cuts, circular order and identifications on a disk. Its principal source is Section 9, "Puzzle and pinched disk models," of Mikhail Lyubich's *Conformal Geometry and Dynamics of Quadratic Polynomials*, in the supplied manuscript dated September 2023. The selected source occupies printed pages 152–163. Section 10, the potential-theory appendix, is not part of this volume's source population.

There are four connected parts. The opening chapters develop compactness, fullness, small paths and quantitative boundary coordinates. The next part examines cut cells, shrinking nests, interior bodies and legal trees. The third constructs closed quotient models and records exactly what a planar recognition theorem needs. The final part develops refinement guarantees, finite noncrossing partitions, electrical and thermal tree models, and the information costs of addresses. Every authored investigation has a worked solution. The source guide retains the original subsection order and all 30 numbered units, including the cases whose literal formulations need repair.

Mathematical scope

The treatment is an expository mathematical development, with detailed written proofs and explicit classical inputs. It is not a claim to have solved a new open problem in complex dynamics. A hull, its physical boundary, an ideal boundary, and a quotient model are separate spaces throughout. Statements about one of them require a proved map before they can be used as statements about another.

The source's broad cut-curve definitions require care. A Jordan cut-curve can intersect a U-shaped arc in a disconnected pair of segments inside its disk. Pairwise nonseparation by cut-lines need not be transitive when an endpoint lies on the separating line: two tangent filled disks give a concrete example. Accordingly, the valid cut-cell and shrinking-neighborhood theorems are proved under stated connectedness and admissibility hypotheses. They are not used to justify a quotient whose proposed fibers overlap. The source guide also distinguishes rigidity of a boundary from rigidity of a filled body, and geometric symmetry from symmetry of chosen internal centers.

For a locally connected hull, the boundary map supplies a genuine closed equivalence relation. Its unlinked classes, convex hulls and quotient receive a separate development. The general non-locally-connected construction in source Proposition 9.29 invokes "cleaning" a lamination without specifying all operations needed for a proof. That general construction remains an explicit obligation. Well-posed conditional quotient statements and concrete examples are proved without declaring that obligation complete. In particular, a universal monotone locally connected quotient and a particular unspecified lamination construction are not identified by terminology alone.

Dependencies and evidence

The undergraduate prerequisite boundary consists of compact metric spaces, elementary real and complex analysis, and finite-dimensional linear algebra. The normalized Riemann map and the Carathéodory boundary theorem are reused from the preceding PHI developments. Jordan separation, Schoenflies extension, arc extraction and Moore's theorem are declared classical topological inputs; their uses and hypotheses are recorded in the dependency register. Other metric and finite-model consequences are developed in the text. The gate, null-family and piecewise-continuity arguments are arranged so that their forward references do not create a circular proof.

Mathematical arguments and source qualifications receive independent agent cross-review. This is distinct from external peer review and from checking in a proof assistant. New Lean, Rocq and CAS development follows manuscript review under the current series instruction. The attached source archive contains editable text, original figure programs and production checks; its checksums and build tests establish artifact identity and reproducibility, not a theorem's truth.

The physical applications are deductions in stated finite network or rod models. Conductances, units, boundary conditions and uncertainty are specified there. No measured dataset, hardware benchmark or experimental validation is asserted. A tree discretization is connected to a continuum only when its representation and error hypotheses are supplied.

Reading presentations

The public and internal editions share the same mathematical core. The internal edition adds teaching routes and assessment material; it does not claim a complete institutional curriculum mapping. Light and dark PDFs contain identical selectable text and navigation, with monochrome vector figures and monospace labels. Dark is the primary reading presentation. One current manuscript is maintained, with stable names and no revision badge.

Contents

Preface and reading contract **i**

1 Compact geometry and the meaning of a model **1**
1.1 Distance to a compact set . . . 1
1.2 Nested sets and uniform trapping . . . 3
1.3 Worked compactness investigations . . . 4

2 Planar hulls and exterior coordinates **6**
2.1 Filling bounded complementary regions . . . 6
2.2 The exterior Riemann coordinate . . . 7
2.3 Worked coordinate investigations . . . 9

3 Local connectivity, small paths and vanishing bodies **10**
3.1 Pointwise and uniform hypotheses . . . 10
3.2 Constructing small paths . . . 11
3.3 Gate maps and a null-family estimate . . . 12
3.4 Worked local-topology investigations . . . 14

4 Circular order, boundary moduli and collar geometry **16**
4.1 Angles and oriented intervals . . . 16
4.2 Moduli inherited from a boundary parametrization . . . 17
4.3 A nonsingular collar and what it controls . . . 18
4.4 Worked order and geometry investigations . . . 18

5 Cut-lines, sectors, and finite cutting **20**
5.1 What a cut is required to do . . . 20
5.2 Several disjoint simple cut-lines . . . 22
5.3 A Jordan cut-curve need not define one cell . . . 23
5.4 Solved investigations . . . 25

6 Countable cutting and the failure of a naive partition **26**
6.1 Retained hulls under infinitely many cuts . . . 26
6.2 Accesses which survive visible-root cutting . . . 27
6.3 Nested compactness gives no rate . . . 29
6.4 Why nonseparation need not be transitive . . . 29
6.5 A valid conditional projection statement . . . 31
6.6 Solved investigations . . . 31

7 Rigidity, local connectivity, and landing **33**
7.1 The topology supplied by a shrinking nest . . . 33
7.2 Finite barriers in the prime-end disk . . . 34
7.3 A local-connectivity criterion using finite sectors . . . 37
7.4 Solved investigations . . . 38

8 Gates and Jordan interior components **39**
8.1 A gate onto one interior component . . . 39
8.2 The boundary of an interior component . . . 42
8.3 The null-family consequence . . . 43

8.4 Solved investigations . . . 43

9 Dipoles, attachments, and peripheral geometry 45
9.1 Dipoles cut a boundary differently from a hull . . . 45
9.2 A precise null-attachment model . . . 46
9.3 Accessibility and the role of the boundary continuum . . . 48
9.4 Dense roots and a null-family criterion . . . 48
9.5 Solved investigations . . . 49

10 Regulated legal arcs and finite trees 51
10.1 Pointing the interior components . . . 51
10.2 Boundary points in the interior of a legal arc . . . 53
10.3 Symmetry includes the centers . . . 55
10.4 The finite legal hull . . . 56
10.5 Solved investigations . . . 57

11 Closed identifications and the topology of a quotient 59
11.1 Closedness and upper semicontinuity . . . 59
11.2 Maps defined by classes . . . 61
11.3 Connected fibers and connected images . . . 62
11.4 Worked investigations on quotient topology . . . 64

12 Unlinked circle fibers and hyperbolic convex hulls 65
12.1 The Klein coordinate and finite convex combinations . . . 65
12.2 Circular order is the separation criterion . . . 68
12.3 Why conformal boundary fibers are unlinked . . . 69
12.4 Worked investigations on hull geometry . . . 69

13 A precise pinched-disk construction 71
13.1 Closedness of the filled relation . . . 71
13.2 The external sphere theorem . . . 74
13.3 Worked investigations on the disk model . . . 75

14 Recognizing a locally connected conformal hull 78
14.1 The boundary correspondence determines the outside . . . 78
14.2 Matching and filling the Jordan regions . . . 79
14.3 Continuity across infinitely many regions . . . 80
14.4 Changing the angular marking . . . 82
14.5 Solved recognition investigations . . . 82

15 Limits and the general model problem 85
15.1 A limit of relations need not be a relation . . . 85
15.2 A precise conditional projection theorem . . . 87
15.3 A Cantor fan and its connected fibers . . . 88
15.4 Locally connected finite fans can converge to this fan . . . 90
15.5 Solved investigations on limits and quotients . . . 91

16 Quantitative refinement and geometric stopping rules 94
16.1 A refinement tree with a certified mesh . . . 94
16.2 Explicit moduli and address metrics . . . 96
16.3 Adaptive refinement and what topology still needs . . . 99
16.4 Solved investigations . . . 100

17 Exact circular order and finite noncrossing certificates 102
17.1 The four-endpoint test . . . 102
17.2 Rational input and independently readable witnesses . . . 103
17.3 A faster certificate for a perfect matching . . . 104
17.4 Uncertain angles and finite evidence . . . 105
17.5 Solved investigations . . . 106

18 Dirichlet problems and elimination on finite trees **108**
18.1 Energy and the unique harmonic extension . . . 108
18.2 Eliminating one vertex exactly . . . 110
18.3 Bottom-up messages on a rooted tree . . . 112
18.4 Solved investigations . . . 113

19 Resistance bounds, perturbations and finite error certificates **115**
19.1 Voltage and current formulations . . . 115
19.2 Monotonicity and uncertainty in conductances . . . 117
19.3 A residual bound with an explicit stability constant . . . 119
19.4 A complete branching model . . . 120
19.5 Solved investigations . . . 120

20 Finite addresses, precision costs and physical interpretation **123**
20.1 Legal addresses recognized by finite states . . . 123
20.2 Depth, storage and geometric accuracy . . . 124
20.3 An exact one-dimensional conduction model . . . 126
20.4 What transfers from a finite model to a continuum . . . 129
20.5 Solved investigations . . . 130

21 Source guide and precise obligations **132**
21.1 9.1. Cut-curves and puzzle pieces for hulls . . . 132
21.2 9.2. Interior components of hulls . . . 135
21.3 9.3. Legal issues . . . 136
21.4 9.4. Pinched disk model for a hull . . . 136

References **138**

Chapter 1

Compact geometry and the meaning of a model

A planar continuum can be described in at least three ways: as a subset of the plane, as the boundary seen from its complement, and as a quotient of a simpler space. These descriptions contain different information. Two external approaches can end at the same point. Two closed pieces can meet at a vertex without belonging to a partition. A sequence of finite drawings can converge to a set with a very different local topology. The purpose of the preliminary chapters is to make the passage between these descriptions precise before constructing a puzzle model.

Our ambient plane is $\mathbb{C} = \mathbb{R}^2$, with Euclidean distance $|z - w|$. A *continuum* is a nonempty compact connected metric space. A planar continuum is allowed to be a singleton unless a statement explicitly excludes that case. The source uses the word hull for a nondegenerate full planar continuum. We normally permit the degenerate case and say *nondegenerate hull* when an exterior Riemann coordinate is required. This convention matters for nested intersections and for the Cantor-fan quotient at the end of the book.

The elementary prerequisites are compactness and continuity in metric spaces, undergraduate complex analysis, and finite-dimensional linear algebra. The Riemann mapping and Carathéodory boundary theorems developed in Volumes VII and VIII are explicitly reused below. Jordan separation, the Schoenflies extension theorem, arc extraction from a path, and the Moore decomposition theorem are named classical topological inputs at their actual points of use. We do not silently treat those results as consequences of a picture. The metric, algebraic, and quantitative consequences needed here receive proofs.

1.1 Distance to a compact set

For a nonempty subset A of a metric space (X, d) put

$$d(x, A) = \inf_{a \in A} d(x, a), \qquad \operatorname{diam} A = \sup_{a,b \in A} d(a, b), \qquad A_\varepsilon = \{x : d(x, A) < \varepsilon\}.$$

The infimum is a minimum when A is compact: choose a minimizing sequence, take a convergent subsequence, and use continuity of distance. In particular, for compact A one has $d(x, A) = 0$ exactly when $x \in A$. If A, B are nonempty compact sets, their Hausdorff distance is

$$d_H(A, B) = \max\left\{\sup_{a \in A} d(a, B), \sup_{b \in B} d(b, A)\right\}.$$

Neither the number of components nor the number of complementary regions is part of this distance. It measures geometric proximity of sets, not agreement of their topology.

The two directed errors can be kept separate throughout a computation:

$$h(A \mid B) = \sup_{a\in A} \min_{b\in B} d(a,b), \qquad d_H(A,B) = \max\{h(A \mid B), h(B \mid A)\}.$$

For $r \geq 0$, compactness gives the exact equivalence

$$h(A \mid B) \leq r \quad \Longleftrightarrow \quad (\forall a \in A)(\exists b \in B)\ d(a,b) \leq r.$$

The order of these quantifiers is essential: the chosen b may depend on a. For $A = [0,1]$ and $B = \{0\}$, the reverse directed error is zero but $h(A \mid B) = 1$. Thus an enclosure and a set approximation are different objects even before questions of topology arise. The zero-error case also explains definiteness: $d_H(A,B) = 0$ implies both inclusions because compact sets are closed. Symmetry, this definiteness, and the triangle estimate below make d_H a metric on the nonempty compact subsets.

Lemma 1.1 (Distance and diameter estimates)**.** *For nonempty compact subsets A, B of a metric space,*

$$|d(x,A)-d(y,A)| \leq d(x,y), \quad |d(x,A)-d(x,B)| \leq d_H(A,B), \quad |\operatorname{diam} A-\operatorname{diam} B| \leq 2d_H(A,B).$$

For three such sets, $d_H(A,C) \leq d_H(A,B) + d_H(B,C)$.

Proof. For every $a \in A$, $d(x,a) \leq d(x,y) + d(y,a)$. Taking infima and interchanging x, y proves the first bound. Write $h = d_H(A,B)$. Given $a \in A$, choose $b \in B$ with $d(a,b) \leq h$. Then $d(x,B) \leq d(x,a) + h$; minimizing over a and interchanging the two sets proves the second bound. For $a, a' \in A$ choose $b, b' \in B$ within h. The triangle inequality gives

$$d(a,a') \leq d(a,b) + d(b,b') + d(b',a') \leq 2h + \operatorname{diam} B.$$

Taking suprema and reversing the sets proves the diameter bound. Finally, choose $b \in B$ near any $a \in A$, then $c \in C$ near b. Taking the two directed suprema yields the Hausdorff triangle inequality. All statements remain valid for singletons.

To spell out the directed triangle step, set $h = d_H(A,B)$ and $k = d_H(B,C)$. For each $a \in A$ choose minimizers $b \in B$ and $c \in C$ successively. Then

$$d(a,C) \leq d(a,c) \leq d(a,b) + d(b,c) \leq h + k.$$

Taking the supremum gives $h(A \mid C) \leq h + k$; interchanging A, C gives the other directed bound. No common nearest-point map or continuity of nearest-point choices was used. Such a map may fail to be unique, while the set-distance estimate remains valid. □

The factor two in the diameter estimate cannot be improved for general compact sets. For $A = [0,1]$ and $B = [-h, 1+h]$ in the real line, $d_H(A,B) = h$ whereas $\operatorname{diam} B - \operatorname{diam} A = 2h$. This elementary sharpness example will be useful whenever a sampled drawing is turned into a diameter bound: a one-sided pointwise tolerance must not be read as an identical diameter tolerance.

Proposition 1.2 (Compactness of compact subsets)**.** *If X is a compact metric space, every sequence of nonempty compact subsets $A_n \subset X$ has a subsequence converging in Hausdorff distance to a nonempty compact subset $A \subset X$. If every A_n is connected, so is A.*

Proof. For each integer $k \geq 1$ choose a finite 2^{-k}-net $F_k \subset X$. Record which balls $B(p, 2^{-k})$, $p \in F_k$, meet A_n. There are finitely many such records. Successive subsequences and a diagonal selection produce a subsequence, still denoted A_n, whose record at each fixed scale is eventually

constant. If A_n, A_m have the same record at scale k, every point of one is within 2^{1-k} of a point of the other. Thus this subsequence is Cauchy in d_H.

Here are the diagonal choices explicitly. Let $I_0 = \mathbb{N}$, and recursively choose an infinite $I_k \subset I_{k-1}$ on which the scale-k record is constant. Choose $n_k \in I_k$ with $n_k > n_{k-1}$. For $p, q \geq k$, both indices belong to I_k. If $a \in A_{n_p}$ and $d(a, v) < 2^{-k}$ for a net point v, the common record supplies $b \in A_{n_q} \cap B(v, 2^{-k})$. Consequently

$$d(a, b) < 2^{1-k}, \qquad d_H(A_{n_p}, A_{n_q}) \leq 2^{1-k} \quad (p, q \geq k).$$

The finite record is only a device for extracting a subsequence; no claim is made that it determines the sets exactly. In particular, two sets with the same record can have different connected components.

Define A to be the set of all limits of sequences $a_j \in A_{n_j}$ with $n_j \to \infty$. Compactness of X makes A nonempty. It is closed: if $a^{(k)} \in A$ tends to a, select $n_k > n_{k-1}$ and $a_k \in A_{n_k}$ with $d(a_k, a^{(k)}) < 1/k$. Then $a_k \to a$. To prove convergence, fix $\varepsilon > 0$ and take N so that $d_H(A_n, A_m) < \varepsilon/2$ for $n, m \geq N$. For $a \in A_n$, select points of later sets within $\varepsilon/2$ and pass to a limit; this gives $d(a, A) \leq \varepsilon/2$. Conversely every $a \in A$ is the limit of later points, each within $\varepsilon/2$ of the fixed compact set A_n, so $d(a, A_n) \leq \varepsilon/2$. Consequently $d_H(A_n, A) \leq \varepsilon/2$ for $n \geq N$.

Suppose the limit were disconnected, $A = C \cup D$, with C, D nonempty disjoint relatively closed sets. They are compact; hence $s = d(C, D) > 0$. The neighborhoods $C_{s/3}$ and $D_{s/3}$ are disjoint. For large n, the set A_n lies in their union and meets both, since it approximates every point of A. Their intersections with A_n then separate A_n, a contradiction.

For precision in the last separation step, take n so large that $d_H(A_n, A) < s/4$. Then

$$A_n \subset C_{s/3} \cup D_{s/3}, \quad A_n \cap C_{s/3} \neq \varnothing, \quad A_n \cap D_{s/3} \neq \varnothing.$$

The first inclusion follows by choosing a point of $A = C \cup D$ within $s/4$ of each point of A_n. The two nonempty intersections follow by approximating a selected point of C and of D. Finally the neighborhoods are disjoint, since a common point would imply $d(C, D) < 2s/3$. Their intersections with A_n are therefore disjoint nonempty relatively open subsets covering A_n, exactly the definition of a separation. □

The proof records where compactness enters. In an unbounded ambient plane, the singletons $\{n\}$ have no convergent subsequence in d_H. In a bounded but incomplete space, the candidate limit can leave the space. When this proposition is used for limbs or puzzle pieces, a common compact ambient set must therefore be identified first.

1.2 Nested sets and uniform trapping

Theorem 1.3 (Nested compacta)**.** *Let $A_1 \supset A_2 \supset \cdots$ be nonempty compact subsets of a metric space. Then $A = \bigcap_n A_n$ is nonempty and compact. For every open set $U \supset A$, there is N with $A_n \subset U$ for $n \geq N$. Hence $d_H(A_n, A) \to 0$ and* $\operatorname{diam} A_n \to \operatorname{diam} A$. *If all A_n are connected, A is connected. In particular, the intersection is a singleton if and only if* $\operatorname{diam} A_n \to 0$.

Proof. If the intersection were empty, the sets $A_1 \setminus A_n$ would form an open cover of A_1, with a finite subcover contradicting the nesting. The same finite-intersection argument applied to $A_n \setminus U$, or a convergent sequence of points chosen outside U, proves eventual containment. For Hausdorff convergence take $U = A_\varepsilon$; the other directed distance vanishes because $A \subset A_n$. Lemma 1.1 then proves the diameter convergence. Connectedness follows from Proposition 1.2, or directly by separating the intersection and using eventual containment in two disjoint neighborhoods. A nonempty compact set has diameter zero exactly when all its points coincide.

One may also see the trapping argument directly with its index order. If no N works for an open $U \supset A$, select increasing n_j and $x_j \in A_{n_j} \setminus U$. Compactness of A_1 gives a convergent subsequence $x_j \to x$. For each fixed m, eventually $n_j \geq m$, so closedness gives $x \in A_m$. Thus

$$x \in \bigcap_{m \geq 1} A_m = A, \qquad x \in X \setminus U,$$

a contradiction. This order, first fix m and then let j tend to infinity, avoids the unjustified assertion that every selected x_j already lies in every A_m. Since $A \subset A_n$, the resulting Hausdorff estimate simplifies to the one directed quantity

$$d_H(A_n, A) = \max_{x \in A_n} d(x, A), \qquad 0 \leq \operatorname{diam} A_n - \operatorname{diam} A \leq 2 d_H(A_n, A).$$

The maximum exists because the distance function is continuous on A_n. □

An escaping curve and a nested family involve different quantifiers. If $\gamma : (0, 1] \to X$ tends to a, and a belongs to the interior of every A_n, then

$$\forall n \; \exists t_n > 0 \; \forall t \in (0, t_n) : \quad \gamma(t) \in A_n.$$

There need not be one positive t_* which works for every n. Conversely, if the displayed condition holds and $\operatorname{diam} A_n \to 0$, then $\gamma(t) \to a$ where $\bigcap A_n = \{a\}$. Indeed choose n with $\operatorname{diam} A_n < \varepsilon$ and then use t_n. This elementary order of choices is the analytic content of many statements that a ray is trapped in a shrinking puzzle nest.

Proposition 1.4 (Cluster sets of a bounded approach)**.** *Let $\gamma : (0, 1] \to X$ be continuous with image in a compact metric space. Its terminal cluster set*

$$\omega(\gamma) = \bigcap_{n \geq 1} \overline{\gamma((0, 1/n])}$$

is a continuum. It is a singleton $\{a\}$ precisely when $\lim_{t \downarrow 0} \gamma(t) = a$. If γ is eventually in a closed set F, then $\omega(\gamma) \subset F$.

Proof. Each tail is connected as the continuous image of an interval, and its closure is compact and connected. Apply Theorem 1.3. If the intersection is $\{a\}$, eventual containment of a tail closure in $B(a, \varepsilon)$ is exactly convergence. The converse follows from the definition of a limit. The last assertion follows because a sufficiently late tail, and hence its closure, lies in F.

This set-theoretic definition agrees with the sequential definition of a cluster value. If a belongs to every tail closure, choose

$$0 < t_n \leq 1/n, \qquad d(\gamma(t_n), a) < 1/n.$$

Then $t_n \to 0$ and $\gamma(t_n) \to a$. Conversely, if $s_j \to 0$ and $\gamma(s_j) \to a$, each fixed tail contains all sufficiently late terms, so a lies in its closure. The equivalence uses no monotonicity of the chosen times. If a monotone sequence is wanted, choose the next time less than half the preceding one while applying the same tail closure property. A single convergent sampling sequence establishes one member of the cluster set, not that this set has only one member. □

1.3 Worked compactness investigations

Investigation 1.5 (A nest with a nontrivial limit)**.** For $A_n = [0, 1] \times [-1/n, 1/n]$, compute the intersection, $d_H(A_n, A)$ and the diameter limit. Explain why an arbitrarily small area does not imply an arbitrarily small diameter.

Solution. The intersection is $A = [0,1] \times \{0\}$. Vertical projection gives $d_H(A_n, A) = 1/n$, with equality at the horizontal edges. The extreme corners give

$$\operatorname{area} A_n = 2/n \longrightarrow 0, \qquad \operatorname{diam} A_n = \sqrt{1 + 4/n^2} \longrightarrow 1.$$

Thus area alone cannot provide the shrinking-gaps estimate needed for a piecewise extension. A long thin gap can have very small area and still contain points separated by a fixed distance.

Investigation 1.6 (A limit invisible to a single sample)**.** Let $\gamma(t) = \exp(i/t)$ for $0 < t \leq 1$. Determine its cluster set and compare it with the subsequence $t_n = (2\pi n)^{-1}$.

Solution. For every $\theta \in [0, 2\pi)$, the times $s_n = (2\pi n + \theta)^{-1}$ give $\gamma(s_n) = e^{i\theta}$ and $s_n \downarrow 0$. Hence the cluster set is the full unit circle. Along t_n, however, $\gamma(t_n) = 1$ identically. Landing of a selected subsequence cannot establish landing of the entire approach. The same distinction applies to a numerical trace of an external ray.

Investigation 1.7 (A quantitative unique limit)**.** Suppose A_n are nested compacta and $\operatorname{diam} A_n \leq Cq^n$ for $C > 0$, $0 < q < 1$. For arbitrary $x_n \in A_n$, prove an a posteriori distance bound to the limit and find a depth sufficient for error $\varepsilon \in (0, C)$.

Solution. The intersection is one point a, and $a, x_n \in A_n$ gives $d(x_n, a) \leq Cq^n$. Thus it suffices that

$$n \geq \left\lceil \frac{\log(C/\varepsilon)}{\log(1/q)} \right\rceil.$$

This is a guarantee only if the diameter bound applies to the exact compact set; a few sampled vertices require an additional enclosure. There is no rounding claim in the real-arithmetic formula for the depth.

Investigation 1.8 (Intersection and perturbation)**.** Give compact sets $A_n \to A$ and $B_n \to B$ for which the intersections do not converge to $A \cap B$, even though all limiting sets are nonempty.

Solution. Take $A_n = A = [0,1] \times \{0\}$ and $B_n = [0,1] \times \{1/n\}$. Then $B_n \to A$ but $A_n \cap B_n$ is empty for every n. The limiting intersection is A. An intersection operation therefore needs a separate stability hypothesis. If one insists on nonempty finite intersections, adjoin the common point $(2, 0)$ to all four sets; the finite intersections are then that singleton while the limiting intersection also contains the unit segment.

Investigation 1.9 (Upper limit and a closed constraint)**.** Suppose A_n lie in a compact space and every convergent sequence $a_{n_j} \in A_{n_j}$ has its limit in a closed set F. Prove that $\sup_{a \in A_n} d(a, F) \to 0$. Is the converse valid?

Solution. Failure supplies $\varepsilon > 0$ and points a_{n_j} at distance at least ε from F. Compactness supplies a convergent subsequence; continuity of distance contradicts its required limit in F. Conversely, if the supremum tends to zero and $a_{n_j} \to a$, then $d(a, F) = 0$, hence $a \in F$. Neither direction says that every point of F is approximated. This is upper convergence, not necessarily Hausdorff convergence to F.

Chapter 2

Planar hulls and exterior coordinates

2.1 Filling bounded complementary regions

A compact connected set $K \subset \mathbb{C}$ is *full* if $\mathbb{C} \setminus K$ is connected. Equivalently its complement has no bounded component: every compact set has precisely one unbounded complementary component, because the exterior of a sufficiently large disk is connected. For a compact continuum E write

$$\mathrm{Fill}(E) = \mathbb{C} \setminus U_\infty(E) = E \cup \bigcup \{U : U \text{ is a bounded component of } \mathbb{C} \setminus E\}.$$

Connectedness of E is part of this definition's use as a hull. Filling two isolated points does not join them. We will not extend a connectedness claim to arbitrary compact sets by changing the noun without changing the proof.

Lemma 2.1 (Fullness and filling)**.** *For a planar continuum E,* $\mathrm{Fill}(E)$ *is a compact full continuum. If $E \subset F$ and F is a full compact set, then* $\mathrm{Fill}(E) \subset F$*. Consequently filling is monotone on continua and* $\mathrm{Fill}(\mathrm{Fill}(E)) = \mathrm{Fill}(E)$*.*

Proof. The unbounded complementary component is open, so its complement is closed; outside a large disk every point belongs to that component, so the complement is bounded. If U is a bounded complementary component, then $\partial U \subset E$ and $\partial U \neq \varnothing$. The set $\overline{U}$ is connected and meets E. The union of E with all such closures is connected, proving the assertion. Its complement is $U_\infty(E)$ by construction.

For the minimality statement, every point of $\mathbb{C} \setminus F$ can be joined within this connected open set to a point outside a large disk. Here connected open subsets of the plane are polygonally connected: the points polygonally reachable from one fixed point form both an open and a relatively closed subset. The path avoids E, hence lies in $U_\infty(E)$. Thus $\mathbb{C} \setminus F \subset U_\infty(E)$ and $\mathrm{Fill}(E) \subset F$. Apply this to the full set $\mathrm{Fill}(F)$ for monotonicity, and to a full E for idempotence.

The boundary inclusion used at the start can also be justified locally. For a complementary component U and $x \in \partial U \setminus E$, there would be a disk $B(x, r) \subset \mathbb{C} \setminus E$. This disk meets U, and their union is connected, forcing the whole disk into that same component. It would make x interior to U, a contradiction. Thus $\partial U \subset E$. Moreover

$$E \cup \bigcup_{U \text{ bounded}} \overline{U} = E \cup \bigcup_{U \text{ bounded}} U = \mathrm{Fill}(E).$$

Every set in this union meets the connected anchor E, which proves connectedness even for infinitely many bounded components. To verify that principle, a separation of the union would put E on one side and force every connected $\overline{U}$ meeting it onto the same side, leaving the other side empty. This is why connectedness of the anchor cannot simply be omitted from a filling argument. □

Lemma 2.2 (A Jordan curve in a full set)**.** *If K is full and $\gamma \subset K$ is a Jordan curve, the bounded Jordan domain U enclosed by γ is contained in* $\operatorname{int} K$*.*

Proof. Jordan separation gives two complementary components for γ. If $z \in U \setminus K$, the connected open set $\mathbb{C} \setminus K$ would contain a path from z to the unbounded side. This path would cross $\gamma \subset K$, impossible. Thus $U \subset K$. Since U is open, $U \subset \operatorname{int} K$. □

This observation is the source of a recurring contradiction. If a closed curve in K can be assembled using arcs associated with two supposedly different interior components, its inside cannot be an omitted hole. The inside joins those components. The hypotheses that the curve is simple, lies in K, and actually borders the relevant open pieces must all be checked; a self-intersecting concatenation is not automatically a Jordan curve.

Proposition 2.3 (Nested full continua)**.** *The intersection of a decreasing sequence of nonempty compact full continua is a full continuum, possibly a singleton.*

Proof. Compactness and connectedness follow from Theorem 1.3. If $A = \bigcap A_n$, then

$$\mathbb{C} \setminus A = \bigcup_n (\mathbb{C} \setminus A_n).$$

The sets on the right are connected and increasing, so their union is connected. To see the last implication directly, any hypothetical separation of the union would meet one sufficiently large member on both sides, contradicting that member's connectedness.

In symbols, if the union were $V \cup W$ with a separation and $v \in V$, $w \in W$, choose indices n_v, n_w for their membership in the increasing family. For $m = \max\{n_v, n_w\}$, both points lie in $\mathbb{C} \setminus A_m$. The two sets

$$V \cap (\mathbb{C} \setminus A_m), \qquad W \cap (\mathbb{C} \setminus A_m)$$

would then separate that connected complement. This proof uses increasing complements and decreasing compact sets; a Hausdorff-convergent family has neither ordering automatically. □

Arbitrary intersections of hulls need not be connected. Two overlapping curved hulls can meet in two separated pieces. Nesting is the feature which turns the finite-intersection argument into a connected limit. Likewise, fullness is not closed under arbitrary Hausdorff convergence. A circle with an open arc of length $1/n$ removed is an arc and hence full; the limit is the whole circle, whose bounded complementary disk is a hole. This is a useful obstruction to inferring fullness from a finite resolution polygonal picture.

2.2 The exterior Riemann coordinate

For a nondegenerate hull K, let $\Omega_K = \widehat{\mathbb{C}} \setminus K$. Its exterior coordinate is the normalized conformal bijection

$$\Phi_K : \{w \in \widehat{\mathbb{C}} : |w| > 1\} \longrightarrow \Omega_K, \qquad \Phi_K(w) = c_K w + b_K + O(w^{-1}), \quad c_K > 0.$$

Existence is the exterior form of the Riemann mapping theorem. The topological fact that the complement of a full continuum is simply connected on the sphere, and the Riemann mapping theorem itself, are the classical planar inputs used in Volume VII. For a singleton K, $\Omega_K = \widehat{\mathbb{C}} \setminus K$ has plane conformal type and is excluded from this normalization. The inverse will be denoted Ψ_K.

For $\theta \in \mathbb{R}/\mathbb{Z}$ and $t > 0$ define

$$R_K(\theta, t) = \Phi_K(e^{t+2\pi i\theta}), \qquad G_K(z) = \log|\Psi_K(z)|, \qquad E_K(t) = \{z \in \Omega_K : G_K(z) = t\}.$$

The level t is a logarithmic exterior coordinate; it is not Euclidean distance to K. Along a ray $G_K(R_K(\theta,t)) = t$. Equipotentials $E_K(t)$ are analytic Jordan curves because Φ_K is conformal on a neighborhood of the circle $|w| = e^t$. Nothing here asserts that the ray lands as $t \downarrow 0$.

Proposition 2.4 (Exterior coordinate identities)**.** *For $z \in \mathbb{C} \setminus K$, the function G_K is positive and harmonic, and*

$$|\nabla G_K(z)| = \left|\frac{\Psi_K'(z)}{\Psi_K(z)}\right|, \qquad G_K(z) = \log|z| - \log c_K + O(|z|^{-1}) \quad (z \to \infty).$$

External rays and equipotentials meet orthogonally at every positive level. If $K' = \alpha K + \beta$ with $\alpha \neq 0$, then

$$G_{K'}(\alpha z + \beta) = G_K(z), \qquad c_{K'} = |\alpha| c_K.$$

Proof. Locally choose a branch of $\log \Psi_K$. Its real part is G_K and its complex derivative is Ψ_K'/Ψ_K, proving harmonicity and the gradient identity through the Cauchy–Riemann equations. The inverse Laurent expansion gives $\Psi_K(z) = z/c_K + O(1)$ and hence the asymptotic formula. In the w-plane, radial and circular tangent vectors are orthogonal; multiplication by the nonzero derivative of Φ_K rotates and scales both, preserving their angle.

Write $\alpha = |\alpha| e^{i\eta}$. The map $\widetilde{\Phi}(w) = \alpha \Phi_K(e^{-i\eta} w) + \beta$ has image $\widehat{\mathbb{C}} \setminus K'$ and leading coefficient $|\alpha| c_K > 0$. Uniqueness of the normalization identifies it with $\Phi_{K'}$. Taking inverse moduli proves the two transformation laws.

More explicitly, if $H = \log \Psi_K = u + iv$ on a local branch, then

$$H' = u_x - i u_y, \qquad |H'|^2 = u_x^2 + u_y^2 = |\nabla G_K|^2, \qquad u_{xx} + u_{yy} = 0.$$

Although v can change by an integral multiple of 2π between branches, $u = \log|\Psi_K|$ and these derivative identities agree on overlaps. For the orthogonality claim, the coordinate tangent vectors are

$$\partial_t \Phi_K(e^{t+2\pi i\theta}) = w \Phi_K'(w), \qquad \partial_\theta \Phi_K(e^{t+2\pi i\theta}) = 2\pi i w \Phi_K'(w).$$

Their real Euclidean inner product is zero because $\operatorname{Re}(v\overline{iv}) = 0$, and neither vector vanishes. For the affine change the inverse identity is

$$\Psi_{K'}(\alpha z + \beta) = e^{i\eta} \Psi_K(z).$$

Taking its modulus removes the angular change and gives the displayed Green-level invariance. Finally write $\Psi_K(z) = (z/c_K)(1 + O(1/z))$ near infinity and use $\log|1+\zeta| = O(|\zeta|)$ for small ζ to obtain the asymptotic formula with its actual $O(|z|^{-1})$ remainder. □

The scalar c_K is the logarithmic capacity normalization. We use the coefficient identity above, not an unproved equivalence with every variational definition of capacity. Potential-theoretic interpretations which require an equilibrium measure are separate results. This keeps the coordinate calculations independent of the next source section's potential-theory appendix.

Theorem 2.5 (Boundary coordinate input)**.** *If a nondegenerate hull K is locally connected, its normalized exterior map has a continuous extension to $|w| \geq 1$ on the sphere. The boundary map*

$$\gamma_K : \mathbb{R}/\mathbb{Z} \longrightarrow \partial K, \qquad \gamma_K(\theta) = \Phi_K(e^{2\pi i\theta})$$

is onto. If K is a closed Jordan disk, this map is a homeomorphism of the circle onto ∂K.

This is the Carathéodory–Torhorst boundary input, established for the series in Volume VIII and stated here with its exact role. It does not assert injectivity for an arbitrary locally connected hull. The segment $[-1,1]$ has two external approaches to each interior point, although it is locally connected. Nor does continuity specify a computable modulus from an arbitrary description of K.

2.3 Worked coordinate investigations

Investigation 2.6 (A round disk)**.** For $K = \overline{B(a,r)}$, $r > 0$, compute Φ_K, G_K, the Euclidean distance to K at level t, and the normal derivative of G_K on the boundary.

Solution. Normalization gives $\Phi_K(w) = a + rw$ and $G_K(z) = \log(|z-a|/r)$. On level t, the distance is $r(e^t - 1)$, while the outward radial derivative at $|z-a| = r$ equals $1/r$. In particular,

$$rt \le r(e^t - 1) \le re^T t \qquad (0 \le t \le T).$$

The lower bound follows by integrating $e^s \ge 1$ and the upper bound by integrating $e^s \le e^T$. Even this smooth example has a scale-dependent conversion between logarithmic level and length.

Investigation 2.7 (The segment and a double boundary trace)**.** Verify that $\Phi(w) = (w + w^{-1})/2$ maps $|w| > 1$ conformally onto $\widehat{\mathbb{C}} \setminus [-1,1]$. Determine its boundary fibers.

Solution. If $\Phi(w) = \Phi(v)$, multiplication and factorization give $(w-v)(1 - 1/(wv)) = 0$. Since $|wv| > 1$, $w = v$. Also $\Phi'(w) = (1 - w^{-2})/2 \neq 0$ in the exterior. Solving the quadratic equation gives roots $w = z \pm \sqrt{z^2 - 1}$ with product one. For $z \notin [-1,1]$ neither root lies on the unit circle, since a unit root would yield $z = \cos\vartheta$. Exactly one root lies outside: this is true near infinity, and cannot change along the connected complement of the segment without a unit-circle root. Thus the map is onto. On the boundary,

$$\gamma(\theta) = \cos(2\pi\theta), \qquad \gamma^{-1}(x) = \{\theta, -\theta\} \quad (-1 < x < 1).$$

At $x = 1, -1$ the two representatives coincide modulo one. The exterior coefficient is $c_K = 1/2$. This explicitly separates a boundary quotient from an injective boundary parametrization.

Investigation 2.8 (A cut cannot pass through an open body)**.** Suppose a continuous injective line $L : \mathbb{R} \to \mathbb{C}$ meets a compact set K in finitely many points. Prove that its image avoids $\operatorname{int} K$.

Solution. If $L(t_0)$ were interior, continuity would place $L(t)$ in K for all t in an interval around t_0. Injectivity makes these infinitely many different points, contradicting the finite-intersection assumption. Thus finite cut-lines of a filled hull only encounter its boundary. This observation is essential when a puzzle construction for ∂K is compared with a purported construction for the entire filled set K.

Investigation 2.9 (An interior component is simply connected)**.** Let D be a component of the interior of a full continuum K. Explain why every Jordan curve in D encloses a region contained in D, and deduce simple connectivity using the planar Jordan criterion.

Solution. Lemma 2.2 puts the bounded region U in $\operatorname{int} K$. Near any point of the Jordan curve, U meets the open set D; hence the connected open set U belongs to the same interior component. Therefore $U \subset D$. The classical planar criterion states that a domain is simply connected precisely when it contains the inside of every Jordan curve it contains. Applying that criterion gives simple connectivity. The criterion is a topological input; it is not being inferred from one sampled family of curves.

Chapter 3

Local connectivity, small paths and vanishing bodies

3.1 Pointwise and uniform hypotheses

A space X is locally connected at x if every neighborhood of x contains a connected relatively open neighborhood of x. It is weakly locally connected at x if the word open is dropped for the smaller connected neighborhood. A neighborhood here contains a relatively open set containing x; it is not merely a connected set containing x. Every singleton is connected, so omitting that qualification would make the condition vacuous.

Proposition 3.1 (From weak local connectivity everywhere)**.** *If X is weakly locally connected at every point, then X is locally connected. Equivalently, every component of every open set in X is open.*

Proof. Let O be open, C one of its components, and $x \in C$. Choose a connected neighborhood N of x contained in O. Since N meets C, maximality of the connected component gives $N \subset C$. There is an open set V with $x \in V \subset N$, so x is interior to C. This holds for every point of C, making C open. Taking the component containing x inside an arbitrary open neighborhood supplies the connected open neighborhood required for local connectivity. The reverse implication follows immediately from the definition.

The key inclusion is $N \cup C \subset O$: it is connected because $N \cap C$ contains x. Maximality of C then says

$$N \cup C = C, \qquad x \in V \subset N \subset C.$$

For the reverse characterization, assume every component of every open set is open. Given $x \in O$ with O open, its component $C_O(x)$ is connected and open and satisfies $x \in C_O(x) \subset O$. This gives the required neighborhood base directly. Notice that the choice of N in the forward implication is made afresh at each $x \in C$, which is the exact place where the hypothesis at every point is needed. □

The proposition is global in its hypothesis. Weak local connectivity at one selected point alone does not justify the step for all the other points in the component of O. The proof explains why a theorem about shrinking nests at every point can imply local connectivity even when the available nests initially give only closed connected neighborhoods.

Theorem 3.2 (Uniform local connectedness)**.** *For a compact metric space X, local connectivity is equivalent to:*

$$\forall \varepsilon > 0 \ \exists \delta > 0 \ \forall x, y \in X,$$
$$d(x, y) < \delta \implies x, y \text{ belong to a connected set of diameter less than } \varepsilon.$$

Proof. Assume local connectivity. Cover X by connected open sets of diameter less than ε. A finite subcover has a Lebesgue number $\delta > 0$: every subset of diameter less than δ lies in one member. For completeness, if no such number existed, finite sets of diameter tending to zero could fail to lie in any member; choose one point from each, pass to a limit x, and use the open covering member containing x to obtain a contradiction. Applying the Lebesgue-number property to $\{x, y\}$ proves the displayed condition.

The finite-cover argument can be given a numerical intermediate step. For each $x \in X$ choose a connected open V_x of diameter less than ε and $r_x > 0$ with $B(x, 2r_x) \cap X \subset V_x$. Choose a finite subcover from the balls $B(x, r_x)$ and set

$$\delta = \min_{1 \le j \le m} r_{x_j} > 0.$$

If $u \in B(x_j, r_{x_j})$ and $d(u, v) < \delta$, then $d(v, x_j) < 2r_{x_j}$, so $u, v \in V_{x_j}$. This proves a uniform input tolerance directly. It does not make the tolerance computable from an arbitrary name for X: obtaining the cover and the radii is itself mathematical input.

Conversely fix an open O and a point x of one component C. Choose $\varepsilon > 0$ with $B(x, 2\varepsilon) \cap X \subset O$. For every $y \in B(x, \delta) \cap X$, where δ is also less than ε, the asserted connected set joins x to y and lies in $B(x, \varepsilon) \cap X \subset O$. Thus $B(x, \delta) \cap X \subset C$. Every component of O is open, and Proposition 3.1 gives local connectivity. □

3.2 Constructing small paths

Theorem 3.3 (Uniform small paths)**.** *If X is compact and locally connected, then for every $\varepsilon > 0$ there is $\delta > 0$ such that any two points at distance less than δ can be joined by a continuous path in X whose image has diameter less than ε. Distinct endpoints can be joined by an arc with the same bound.*

Proof. We give the path construction to specify the uniformity. Choose positive numbers r_k with $2\sum_{k \ge 0} r_k < \varepsilon$; for example $r_k = \varepsilon 2^{-k-4}$. By Theorem 3.2, choose $\delta_k > 0$ so that points within δ_k belong to a connected set of diameter less than r_k.

Whenever two points belong to a connected set C, they can be joined by a finite η-chain of points of C for every $\eta > 0$. Indeed the chain-reachable subset from the first point and its complement are both relatively open in C; connectedness makes the reachable subset all of C. Start with x, y at distance less than δ_0 and a connected carrier C_0 of diameter less than r_0. Choose a finite δ_1-chain from x to y inside C_0, and assign its vertices in order to a partition of $[0, 1]$. Insert repeated vertices if needed so every interval length is at most $1/2$.

For the chain assertion, define $u \equiv_\eta v$ when some finite sequence $u = u_0, u_1, \ldots, u_m = v$ in C has $d(u_{j-1}, u_j) < \eta$ for every j. Concatenation and reversal show that this is an equivalence relation. If v is in one class, then $B(v, \eta) \cap C$ is in that class too; hence each class is relatively open, and its complement is a union of other open classes. Connectedness leaves exactly one class. The finite sequence need not have a length bounded independently of C or η.

For each adjacent pair of level-one vertices choose a connected carrier of diameter less than r_1 containing them. Subdivide its parameter interval using a finite δ_2-chain in that carrier, with all new interval lengths at most half the parent's. Continue inductively. Values already assigned at endpoints are retained, all partition meshes tend to zero, and the assigned points lie in X. Descendants of any level-k interval lie within distance $\sum_{j \ge k} r_j$ of either one of its endpoint values: the successive carriers add at most the stated tail. Thus the assignment on the dense union of all partition vertices is uniformly continuous. More explicitly, fix k. Sufficiently close parameters belong to the same or adjacent level-k intervals; their values differ by at most $2\sum_{j \ge k} r_j$ through a common endpoint. The tail tends to zero.

To make the parameter estimate precise, let P_k be the finite partition at stage k and let $\mu_k > 0$ be the smallest length of its intervals. Two parameters at distance less than μ_k lie in the same interval or in adjacent intervals, including their endpoints. With

$$T_k = \sum_{j \geq k} r_j = \varepsilon 2^{-k-3} \quad \text{for the chosen geometric sequence,}$$

the values assigned to parameters in adjacent intervals differ by at most $2T_k$ through their common endpoint. Given an output tolerance $a > 0$, choose k with $2T_k < a$ and then use input tolerance μ_k. The order is crucial: no positive lower bound on all μ_k is needed. The partition meshes tend to zero, making the assigned parameter set dense, while the tail estimate makes its image assignment uniformly continuous. These are two distinct ingredients in the extension step.

Completeness of compact X extends this assignment uniquely to a continuous map $\alpha : [0,1] \to X$ with endpoints x, y. Every point of the path is within $\sum_{j \geq 0} r_j$ of x, so its diameter is less than ε. The classical arc-extraction lemma for a path in a Hausdorff space gives an injective subpath image joining distinct endpoints; its image is contained in $\alpha([0,1])$ and has no larger diameter. This last loop-removal lemma is the named topological input in the arc assertion; the quantitative path construction itself has been given above.

For completeness, if D_* denotes the dense set of assigned parameters and $t \in [0,1]$, choose $s_n \in D_*$ with $s_n \to t$. Uniform continuity makes the sequence of assigned values Cauchy; completeness gives its limit. If another sequence $s'_n \to t$ is used, then $|s_n - s'_n| \to 0$, so the two image sequences have the same limit. This defines $\alpha(t)$ independently of choices. The same uniform estimate passes to the limits and proves continuity. Finally

$$d(\alpha(t), x) \leq T_0, \qquad \operatorname{diam} \alpha([0,1]) \leq 2T_0 = \varepsilon/4 < \varepsilon.$$

The spare factor is deliberate; no limiting equality can spoil the strict diameter target. Arc extraction only shrinks the image set, so this bound survives removal of loops. □

A connected compact locally connected metric space is therefore arcwise connected. One may prove the global assertion without making a new uniform estimate: the set of points path-reachable from a fixed point is open by the local path property, and its complement is open by the same argument. Connectedness makes the complement empty. Arc extraction then replaces a path by an arc whenever its endpoints differ.

3.3 Gate maps and a null-family estimate

The topology chapter proves a specific gate theorem (Theorem 8.1): for a locally connected full continuum K and a component D of $\operatorname{int} K$, there is a retraction $\pi_D : K \to \overline{D}$. It is constant on any arc disjoint from $\overline{D}$; if an arc meets $\overline{D}$, its gate image is contained in that arc. It is locally constant on $K \setminus \overline{D}$. The proof uses small arcs and Jordan fullness, not the null-family conclusion below. The following forward reference is therefore acyclic: small paths imply the gate theorem, which then implies the shrinking-body theorem.

Proposition 3.4 (A common modulus for gate maps)**.** *For each $\varepsilon > 0$, choose δ from Theorem 3.3 for the compact locally connected hull K. Then, for every interior component D and all $x, y \in K$,*

$$|x - y| < \delta \quad \Longrightarrow \quad |\pi_D(x) - \pi_D(y)| < \varepsilon.$$

The same δ works for every component.

Proof. Join x to y by an arc A of diameter less than ε. If A misses $\overline{D}$, its gate image is a singleton. Otherwise $\pi_D(A) \subset A$. In both cases the displayed distance is less than ε. For $x = y$ no arc is needed. The selected small-arc bound depends only on K, not on D.

Equivalently, when interior components exist, define a family oscillation at scale s by

$$W(s) = \sup_{D} \sup_{\substack{x,y\in K \\ |x-y|\le s}} |\pi_D(x) - \pi_D(y)|.$$

If there are no interior components, set $W(s) = 0$. The proposition gives $W(s) \to 0$ as $s \downarrow 0$: for the strict input bound δ take $s \le \delta/2$. The supremum over D is taken before the limit in s. Exchanging this order would give only individual continuity and would not control a moving sequence of components. If there are no interior components the family assertions are vacuous; the null-family conclusion then holds immediately. □

Theorem 3.5 (Interior components form a null family)**.** *Let $K \subset \mathbb{C}$ be a compact locally connected full continuum. For every $\varepsilon > 0$ there are only finitely many components D of* int K *with* diam $D \ge \varepsilon$*. Equivalently, the diameters of any sequence of pairwise distinct interior components tend to zero.*

Proof. Suppose distinct components D_n have diameters at least $\varepsilon > 0$. Select $x_n, y_n \in D_n$ with $|x_n - y_n| > \varepsilon/2$. By compactness, pass to subsequences with $x_n \to x$, $y_n \to y$, and $|x - y| \ge \varepsilon/2$.

If $n \ne m$, the connected open set D_n is disjoint from $\overline{D_m}$. Indeed a point of $D_n \cap \overline{D_m}$ has a small disk contained in D_n which also meets D_m, contradicting distinctness of the components. Local constancy of π_{D_m} outside $\overline{D_m}$ therefore gives

$$\pi_{D_m}(x_n) = \pi_{D_m}(y_n) = a_{m,n} \qquad (n \ne m).$$

Choose the common gate modulus with output tolerance $\varepsilon/8$. For distinct sufficiently large m, n, both $|x_m - x_n|$ and $|y_m - y_n|$ are below its input tolerance. Because π_{D_m} fixes x_m, y_m,

$$|x_m - y_m| \le |\pi_{D_m}(x_m) - a_{m,n}| + |a_{m,n} - \pi_{D_m}(y_m)| < \varepsilon/4,$$

contradicting the choice of the points. Countability of the family, when needed, follows independently because disjoint nonempty open subsets of the plane contain different points of $\mathbb{Q}^2$.

The quantifiers behind the contradiction are worth isolating. Uniform continuity supplies one $\delta > 0$ for $\varepsilon/8$. Since both endpoint sequences converge, there is one N such that

$$m, n \ge N \quad \Longrightarrow \quad |x_m - x_n| < \delta, \quad |y_m - y_n| < \delta.$$

Now choose any two distinct indices $m, n \ge N$. Retraction, constancy on the other component, and the common modulus give in that order

$$|x_m - y_m| \le W(|x_m - x_n|) + W(|y_m - y_n|) \le \varepsilon/4.$$

One may use the proposition directly instead of W when a strict supremum bound is not attained. The contradiction does not require choosing a gate which is continuous at an unknown limiting point after its index has already varied. For countability, enumerate $\mathbb{Q}^2$ and associate to each component the first rational point it contains; disjoint open components receive distinct indices. □

The conclusion uses local connectivity of the fixed hull. It is stronger than saying the areas of disjoint bounded domains tend to zero: area summability holds without local connectivity, while long narrow domains can retain a positive diameter. The proof also exposes why pointwise continuity of each gate would be insufficient. The moving component D_m requires one modulus for the entire family of gates.

Lemma 3.6 (Pasting across a null family)**.** *Let X be compact metric, Y a metric space, $A \subset X$ closed, and let U_j be pairwise disjoint open sets with $X = A \cup \bigcup_j U_j$ and $\partial_X U_j \subset A$. Suppose $f : X \to Y$ is continuous on A and on every $\overline{U_j}$, and its restrictions agree on overlaps. Assume every $x \in U_j$ has $d(x, \partial_X U_j) \leq \eta_j$, where $\eta_j \to 0$ along distinct indices, and* $\operatorname{diam} f(\overline{U_j}) \to 0$ *along distinct indices. Then f is continuous on X.*

Proof. Take $x_n \to x$. If infinitely many x_n lie in A, that subsequence is controlled by continuity on A. If infinitely many lie in one fixed U_j, the limit lies in $\overline{U_j}$ and continuity there applies. The remaining case has $x_n \in U_{j_n}$ with distinct indices. Choose $a_n \in \partial_X U_{j_n}$ with $d(x_n, a_n) \leq \eta_{j_n} + 1/n$. Then $a_n \to x$, so $x \in A$, and

$$d_Y(f(x_n), f(x)) \leq \operatorname{diam} f(\overline{U_{j_n}}) + d_Y(f(a_n), f(x)) \longrightarrow 0.$$

Every subsequence has been covered by these alternatives; otherwise a subsequence staying a positive distance from $f(x)$ would contradict one of them. This proves continuity.

Here is the exhaustive subsequence reduction used in the last sentence. If continuity failed at x, there would be $b > 0$ and a sequence $x_n \to x$ with $d_Y(f(x_n), f(x)) \geq b$. Either infinitely many terms lie in A, or discard the finitely many which do. In the second case, either one index occurs infinitely often or a subsequence has all indices distinct. The first two cases contradict continuity on a fixed closed piece. In the distinct-index case,

$$d_X(a_n, x) \leq \eta_{j_n} + 1/n + d_X(x_n, x) \longrightarrow 0,$$

and closedness of A puts x in A. Given b, choose n large enough that both the target diameter and $d_Y(f(a_n), f(x))$ are less than $b/3$. The displayed inequality in the proof is then less than $2b/3$, a contradiction. Agreement on overlaps ensures that $f(a_n)$ is the same value whether evaluated through A or through $\overline{U_{j_n}}$. □

3.4 Worked local-topology investigations

Investigation 3.7 (The comb obstruction)**.** Let

$$C = ([0,1] \times \{0\}) \cup (\{0\} \times [0,1]) \cup \bigcup_{n \geq 1} (\{1/n\} \times [0,1]).$$

Prove compactness and connectedness, and show that C is not locally connected at $(0, t)$ when $0 < t \leq 1$.

Solution. Every convergent sequence in the displayed union either stays on a fixed tooth or has tooth coordinates tending to zero, in which case its limit lies on the included vertical segment. Thus the bounded set is closed and compact. Each vertical segment meets the horizontal base, so the union is connected. Fix $0 < r < t/3$. In $C \cap B((0,t), r)$ the connected component containing $(0, t)$ lies on the vertical line $x = 0$: projection to the first coordinate maps a connected set to an interval, while the available coordinates form $\{0\} \cup \{1/n\}$, containing no nontrivial interval. Every relative neighborhood of $(0, t)$ nevertheless meets other teeth. Therefore that component is not a neighborhood, so no connected open neighborhood can lie in the ball.

Investigation 3.8 (Nonuniform small-scale information)**.** Explain why the existence of a small-path modulus for each compact locally connected set does not supply one common modulus for all polygonal arcs of diameter at most two.

Solution. Take a U-shaped polygonal arc whose two upper endpoints are distance h apart and whose vertical sides have height one. The only arc in that set joining the endpoints travels

down to the base and back up, so its diameter is at least one. Letting $h \downarrow 0$ defeats every proposed common positive input tolerance for output diameter $1/2$. Each fixed polygonal arc still has its own modulus. Uniformity over a family is extra data, not a consequence of pointwise compactness.

Investigation 3.9 (Why a boundary-distance condition suffices)**.** Let U_j be disjoint bounded open subsets of a common bounded planar region. Show that $\sup_{z \in U_j} d(z, \partial U_j) \to 0$ along distinct indices, even if their diameters do not tend to zero.

Solution. If the supremum exceeded $r > 0$ for infinitely many j, each such U_j would contain a disk of radius $r/2$. These disjoint disks would have infinite total area in a bounded region. Quantitatively,

$$\#\{j : \sup_{z \in U_j} d(z, \partial U_j) > r\} \leq \frac{4 \operatorname{area}(B)}{\pi r^2}$$

when all the sets lie in a measurable bounded region B. This controls the source-side boundary approximation in Lemma 3.6. It supplies no target-side diameter bound; that is a separate hypothesis.

Investigation 3.10 (A common-modulus contradiction)**.** Abstract the proof of Theorem 3.5: let disjoint sets E_n in a compact metric space admit retractions r_n fixing E_n, constant on every E_m with $m \neq n$, and equicontinuous as a family. Prove $\operatorname{diam} E_n \to 0$ without using planarity.

Solution. If not, choose pairs $x_n, y_n \in E_n$ separated by a fixed positive amount and pass to convergent subsequences. Uniform continuity of all r_m makes $r_m(x_n)$ close to x_m and $r_m(y_n)$ close to y_m when m, n are large. Constancy on E_n equates these two images, contradicting their separation. Planarity is used earlier to construct the gates, not in this final compactness argument. This separation of roles prevents the abstract lemma from being mistaken for a theorem that arbitrary compact spaces possess gate retractions.

Investigation 3.11 (A filled set and its boundary)**.** Compare local connectivity and finite-cut rigidity for $K = \overline{\mathbb{D}}$ and $J = \partial\mathbb{D}$.

Solution. Both sets are locally connected. Any injective curve intersecting K only finitely many times avoids its interior by Investigation 2.8; it cannot split the open disk into two pieces. Small cuts through the complementary interior of J, however, can isolate short circular arcs of J. Thus a boundary puzzle which uses both complementary components of J is not a hull puzzle for K. Local connectivity by itself does not erase the difference between the two admissible-cut classes.

Chapter 4

Circular order, boundary moduli and collar geometry

4.1 Angles and oriented intervals

Write $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, with $e(\theta) = e^{2\pi i\theta}$. An ordered triple of distinct angles (a, b, c) is positively cyclically ordered if one can choose lifts satisfying $\widetilde{a} < \widetilde{b} < \widetilde{c} < \widetilde{a} + 1$. The interval $(a, b)_{\mathbb{T}}$ is the positively oriented arc from a to b, excluding both endpoints. Its length lies in $(0, 1)$ when $a \neq b$. We never choose a real branch of an angle without recording where the circle has been cut.

The circular metric is

$$d_{\mathbb{T}}(a, b) = \min_{k\in\mathbb{Z}} |\widetilde{a} - \widetilde{b} - k|, \qquad 0 \leq d_{\mathbb{T}}(a, b) \leq \tfrac{1}{2}.$$

For $s = d_{\mathbb{T}}(a, b)$, elementary trigonometry yields

$$|e(a) - e(b)| = 2\sin(\pi s), \qquad 4s \leq |e(a) - e(b)| \leq 2\pi s.$$

The lower bound follows from concavity of sine on $[0, \pi/2]$ and the chord joining its endpoint values; the upper bound follows from $\sin u \leq u$. These estimates allow conversion between angular and Euclidean tolerances on the circle, with explicit constants.

Lemma 4.1 (Alternating endpoints)**.** *Let a, b, c, d be four distinct points of the unit circle. The open straight chords (a, b) and (c, d) intersect if and only if their endpoints alternate in circular order. Their intersection, when it exists, is a single point.*

Proof. The line through a, b divides the circle, apart from a, b, into two open arcs lying in opposite open half-planes. The segment from c to d crosses this line precisely when c, d lie in opposite half-planes. In that case its intersection with the line lies in the closed disk, and the intersection of this line with the disk is exactly the chord $[a, b]$. The endpoints are distinct, so the crossing is interior to both segments. Conversely an interior crossing puts c, d on opposite sides, since their chord is not collinear with $[a, b]$: a line meets the circle in at most two points. Opposite sides are precisely the alternating-endpoint condition. Two noncollinear lines meet once.

There is an algebraic form of the same separation argument. For complex vectors define $[u, v] = \mathrm{Im}(\overline{u}v)$ and put $s_c = [b - a, c - a]$, $s_d = [b - a, d - a]$. The distinct-circle-point hypotheses make both nonzero. Opposite sides mean $s_c s_d < 0$, in which case

$$t = \frac{s_c}{s_c - s_d} \in (0, 1), \qquad z = (1 - t)c + td, \qquad [b - a, z - a] = 0.$$

Since c, d are distinct boundary points, their open segment is inside the open disk; hence z lies strictly between a and b on the line's disk intersection. The formula gives an actual intersection

parameter after the order test, but computing it is unnecessary if only the crossing decision is requested. The sign test is strict, so shared endpoints and repeated chords require their separate cases. □

For arbitrary embedded crosscuts of a disk the same separation statement follows from the Jordan theorem: join one crosscut to either boundary arc and inspect the two sides. Straight chords make the condition an exact order test. Hyperbolic geodesics have the same crossing rule, as the Klein-coordinate construction in the pinched-model chapter will prove. The order statement survives even when distance estimates in a particular model do not.

4.2 Moduli inherited from a boundary parametrization

For a continuous map $f : X \to Y$ on a compact metric space define its oscillation modulus

$$\omega_f(r) = \sup\{d_Y(f(x), f(y)) : d_X(x, y) \le r\}, \qquad r \ge 0.$$

It is nondecreasing, finite, and tends to zero as $r \downarrow 0$. Indeed failure of the last assertion produces x_n, y_n with $d(x_n, y_n) \to 0$ but separated images; a convergent subsequence of x_n and then y_n contradicts continuity at their common limit. We do not assert strict increase or invertibility of this modulus.

Proposition 4.2 (A boundary image estimate)**.** *Let $\gamma : \mathbb{T} \to \mathbb{C}$ be continuous, with modulus measured in $d_{\mathbb{T}}$. If a set $A \subset \mathbb{T}$ is contained in an arc of angular length $\ell \le 1/2$, then*

$$\operatorname{diam} \gamma(A) \le \omega_\gamma(\ell).$$

If $\widetilde{\gamma}$ is another map with $\sup_\theta |\gamma(\theta) - \widetilde{\gamma}(\theta)| \le \eta$, then

$$d_H(\gamma(A), \widetilde{\gamma}(A)) \le \eta, \qquad \operatorname{diam} \widetilde{\gamma}(A) \le \omega_\gamma(\ell) + 2\eta$$

for nonempty compact A.

Proof. Any two points of A are at circular distance at most ℓ, so the first estimate is the definition of the modulus. For the second, pair the two image points with the same parameter. The pairwise distance is at most η in both directed Hausdorff estimates. The diameter bound follows either from Lemma 1.1 or the three-term triangle inequality through $\gamma(\theta), \gamma(\phi)$.

The three-term version keeps the two sources of error visible:

$$\begin{aligned} |\widetilde{\gamma}(\theta) - \widetilde{\gamma}(\phi)| &\le |\widetilde{\gamma}(\theta) - \gamma(\theta)| + |\gamma(\theta) - \gamma(\phi)| + |\gamma(\phi) - \widetilde{\gamma}(\phi)| \\ &\le 2\eta + \omega_\gamma(\ell). \end{aligned}$$

Thus refining the angular arc only changes the modulus term; it does not reduce the two endpoint evaluation errors. The same-parameter coupling in the Hausdorff estimate needs no inverse parametrization, so it remains valid when many different angles land at the same physical point. For a nonclosed A the diameter bound still makes sense, but the compact-set Hausdorff convention used in the second part requires the explicitly stated compactness assumption. □

For a locally connected hull, Theorem 2.5 supplies such a map γ_K. It supplies existence of ω_{γ_K}, but not numerical values. A quantitative puzzle estimate which inserts $\omega_{\gamma_K}(2^{-n})$ is a conditional error bound until a modulus is part of the input or independently bounded. The distinction does not weaken the qualitative theorem; it states its effective content.

4.3 A nonsingular collar and what it controls

Fix $1 < r < R < \infty$ and a conformal exterior map Φ_K. It is holomorphic on a neighborhood of the compact annulus $A_{r,R} = \{r \le |w| \le R\}$. Put

$$M_{r,R} = \max_{w \in A_{r,R}} |\Phi_K'(w)|, \qquad m_{r,R} = \min_{w \in A_{r,R}} |\Phi_K'(w)| > 0.$$

Positivity of the minimum follows from injectivity, which forces the complex derivative to be nonzero, and compactness. These are local derivative bounds. They do not automatically give a global lower Lipschitz bound for pairs whose joining segment leaves the annulus.

Proposition 4.3 (A collar cell diameter bound)**.** *Let $r \le r_1 \le r_2 \le R$ and let I be an angular interval of length $\ell \le 1/2$. For $C = \{\rho e(\theta) : r_1 \le \rho \le r_2, \theta \in I\}$,*

$$\operatorname{diam} \Phi_K(C) \le M_{r,R}((r_2 - r_1) + 2\pi R\ell).$$

If a polygonal or sampled enclosure $\widetilde{C}$ of the image has Hausdorff error at most η, its diameter differs from the exact image diameter by at most 2η.

Proof. Join any two points of C by a radial segment followed by a circular arc within C. The radial length is at most $r_2 - r_1$, and the circular length at most $2\pi R\ell$. Along either smooth piece,

$$|\Phi_K(w_1) - \Phi_K(w_0)| \le \int |\Phi_K'(w)|\,|dw| \le M_{r,R}\,\text{length(path)}.$$

Taking the supremum over pairs proves the result. The enclosure claim is the diameter estimate from Lemma 1.1.

Choose one real lift of I and write the two points as $w_0 = \rho_0 e(\theta_0)$, $w_1 = \rho_1 e(\theta_1)$. A concrete joining path has its two pieces

$$s \longmapsto ((1-s)\rho_0 + s\rho_1)e(\theta_0), \qquad s \longmapsto \rho_1 e((1-s)\theta_0 + s\theta_1), \quad 0 \le s \le 1.$$

Their lengths are $|\rho_1 - \rho_0|$ and $2\pi\rho_1|\theta_1 - \theta_0|$. The lift ensures the second path stays in the chosen cell, even when the interval crosses the conventional angular cut at zero. Integrating the derivative on these paths avoids using the straight segment, which could leave the annulus where the bound on Φ_K' was established. The minimum derivative bound also does not reverse the integral triangle inequality: in general

$$\left|\int v(s)\,ds\right| \le \int |v(s)|\,ds$$

has no reverse bound without control of the directions of $v(s)$. This is the precise obstruction to obtaining a global lower distance estimate merely by substituting $m_{r,R}$ for $M_{r,R}$. □

No assertion is made that $M_{r,R}$ stays bounded as $r \downarrow 1$. For domains with cusps, corners or severe boundary distortion, derivatives can behave singularly near the ideal circle. For nonsingular positive levels the estimate is complete; passing to the physical boundary needs the modulus in Proposition 4.2 or an additional regularity theorem. A diagram can depict both regimes without certifying that their constants agree.

4.4 Worked order and geometry investigations

Investigation 4.4 (An exact circular crossing test)**.** Take angles $a = 1/10$, $b = 3/5$, $c = 2/5$, $d = 9/10$. Decide whether the chords $[e(a), e(b)]$ and $[e(c), e(d)]$ cross, without computing sines or cosines. Then replace c by $4/5$.

Solution. Initially $a < c < b < d < a + 1$, so the endpoints alternate and the chords cross. After replacement both $c = 4/5$ and $d = 9/10$ lie on the same oriented arc from b to a, so they do not cross. Rational order comparisons suffice. Equal endpoints would require a separate shared-vertex case and must not be accepted as an interior crossing by a strict test.

Investigation 4.5 (Why the circle has no preferred cut)**.** An arc runs from 9/10 to 1/10 positively around the circle. Compute its length and describe a valid lift for evaluating its order tests.

Solution. Use the lifted interval $[9/10, 11/10]$, of length $1/5$; the endpoint $11/10$ represents $1/10$. The naive difference $1/10 - 9/10 = -4/5$ is not the length of this arc. A consistent algorithm first chooses a basepoint away from the relevant endpoints or uses modular comparisons. Changing the lift adds one to appropriate coordinates without changing cyclic order, provided the same convention is applied to every endpoint.

Investigation 4.6 (A Hölder boundary estimate)**.** Suppose $|\gamma(a) - \gamma(b)| \le H d_{\mathbb{T}}(a, b)^{\alpha}$ with $H > 0$, $0 < \alpha \le 1$. Bound the image diameter of a dyadic arc of length 2^{-n} for $n \ge 1$, and include a uniform map error η.

Solution. Proposition 4.2 gives

$$\operatorname{diam} \gamma(I) \le H 2^{-\alpha n}, \qquad \operatorname{diam} \widetilde{\gamma}(I) \le H 2^{-\alpha n} + 2\eta.$$

To guarantee the latter is at most ε, first require $2\eta < \varepsilon$, then choose $n \ge \alpha^{-1} \log_2(H/(\varepsilon - 2\eta))$, rounded upward and bounded below by one. A refinement of the angular grid cannot overcome a fixed error floor 2η in the map itself.

Investigation 4.7 (Finite contacts do not define a partition)**.** Let $A = [-1, 0]$ and $B = [0, 1]$. Explain why the family $\{A, B\}$ is a cover by compact connected sets but is not the family of fibers of a function. Give two different ways to repair it and state what each repair changes.

Solution. Fibers of a function are disjoint, whereas $A \cap B = \{0\}$. One can replace the cover by the partition $\{[-1, 0), \{0\}, (0, 1]\}$; its nontrivial pieces are no longer compact. Alternatively, impose the equivalence relation generated by declaring each original set a class. Since both contain zero, transitivity merges them into $[-1, 1]$, destroying the original two-piece distinction. Neither repair leaves every requested property unchanged. A quotient construction must specify which mathematical property it intends to retain.

For the second repair, the transitive implication can be written without any picture:

$$x \in A,\ y \in B \quad \Longrightarrow \quad x \sim 0 \sim y \quad \Longrightarrow \quad x \sim y.$$

Thus every equivalence relation containing the identifications internal to A and B contains the universal relation on $[-1, 1]$. The first repair has a different defect: sequences in $[-1, 0)$ can converge to 0 outside that piece. A compact-to-Hausdorff continuous map has closed fibers, since the inverse image of a closed singleton is closed, so this particular noncompact partition cannot be the fibers of such a map.

Investigation 4.8 (Nested full sets versus a converging slit)**.** Let K_n be the unit circle with the open arc of angles $(-1/n, 1/n)$ removed, for $n \ge 3$. Prove that the K_n are full continua converging in Hausdorff distance to a non-full continuum. Why does this not contradict Proposition 2.3?

Solution. Each K_n is a closed circular arc and its complement is connected: the gap joins the inside and outside. The omitted arc has diameter tending to zero, so K_n converges to the whole unit circle. Its complement consists of the open disk and the exterior, so it is not full. The family K_n is increasing rather than decreasing, and its closure of union adds the missing point. The hypotheses of the nested full-continuum theorem are therefore absent. This example will also warn against treating a finite approximation as a topological certificate.

Chapter 5

Cut-lines, sectors, and finite cutting

5.1 What a cut is required to do

Source correspondence. Lyubich, §9.1.1 and Lemma 9.1.

A compact connected subset of the plane will be called a continuum. In this chapter a hull is a nondegenerate continuum whose complement is connected. Nondegenerate means that it contains at least two points. We allow singletons as limits, but name them separately whenever nondegeneracy is an assumption. This distinction matters: an intersection of genuine subhulls can be one point. Fullness is a condition on the entire complement, whereas connectedness is a condition on the set itself. Neither condition may be silently substituted for the other.

Fix a hull $K \subset \mathbb{C}$ and its exterior $\Omega = \mathbb{C} \setminus K$. An access to $a \in \partial K$ is an equivalence class of curves in Ω landing at a, with the local-component interpretation established in Volume VIII. Explicitly, two such curves represent the same access when

$$\forall \varepsilon > 0 \ \exists t_0 > 0 : \quad \gamma_1((0,t_0)), \gamma_2((0,t_0)) \text{ lie in one component of } \Omega \cap B(a,\varepsilon).$$

One may use spherical balls instead. On a fixed bounded neighborhood of a the two metrics are uniformly comparable, so the equivalence relation is unchanged. The definition concerns complete tails, not a selected sequence of points on each curve.

Definition 5.1 (Simple cut-line)**.** A simple cut-line for K is a proper embedding $\ell : \mathbb{R} \to \mathbb{C}$ whose image L meets K at exactly one point $a = \ell(0)$, and whose two local branches at a represent distinct accesses. Its two complementary components are its open sectors S^+ and S^-. The point a is its root.

Properness means that inverse images of compact sets are compact. Thus $L \cup \{\infty\}$ is a Jordan curve in the sphere. The Jordan and Schoenflies theorems give exactly two sectors and a homeomorphism of the plane carrying the pair (L, a) to the real line with marked point 0. The two unbounded ends have a single common endpoint in the sphere; they are not two independent boundary points of a compact interval.

The distinct-access condition excludes a line which merely touches a one-sided hull. To see its role, suppose K does not meet S^+. After straightening the line, arbitrarily small semicircles in the upper half-plane connect its two local branches while avoiding K. Pulling them back gives connections within every prescribed neighborhood of a. The branches then represent the same access. Consequently a cut-line satisfies

$$K \cap S^+ \neq \varnothing, \qquad K \cap S^- \neq \varnothing.$$

A conformal external ray pair landing at one point gives an important example when the two rays represent different accesses. The cut-line consists of their union and the common landing

point. Each ray tends to infinity at its other end; its normalization does not affect the resulting sectors.

Lemma 5.2 (One cut preserves a continuum)**.** *Let K be a hull and L a simple cut-line rooted at a. For either sector S the set*

$$K_S = K \cap \overline{S} = (K \cap S) \cup \{a\}$$

is a nondegenerate hull, and $K_S = \overline{K \cap S}$. The two resulting subhulls meet exactly at a.

Proof. Compactness follows from the first formula. Suppose $K_S = A \sqcup B$ is a separation into nonempty relatively open and closed sets. Both are compact. Exactly one contains a; relabel so that $a \notin B$. Then $B \subset S$. Because B is relatively open in K_S, each of its points has an open plane neighborhood whose intersection with K_S lies in B. Intersect that neighborhood with S. Its intersection with K still lies in B. Therefore B is open in K. It is closed in K because it is compact. This contradicts connectedness of K.

Here is the neighborhood calculation with its quantifiers exposed. For each $b \in B$, relative openness supplies $r_b > 0$ such that $B(b, r_b) \cap K_S \subset B$. Since $b \in S$ and S is open, decrease r_b until $B(b, r_b) \subset S$. Then

$$B(b, r_b) \cap K = B(b, r_b) \cap K_S \subset B, \qquad B = K \cap \bigcup_{b \in B} B(b, r_b).$$

This proves openness in the larger space K, which does not follow from relative openness in K_S alone without the side condition. Compactness also gives $B = K \cap F$ for the closed plane set $F = B$. Thus B is both open and closed in K, is nonempty, and misses a. Those three properties are the precise contradiction; no path connectivity of K was used.

Write S' for the opposite sector. The complement identity

$$\mathbb{C} \setminus K_S = (\mathbb{C} \setminus K) \cup S'$$

proves fullness without an implicit bounded-component argument. Both sets on the right are connected, and their intersection is nonempty because S' is unbounded whereas K is compact. The sector contains a point of K, by the access argument preceding the lemma. Thus K_S has that point and a, so is nondegenerate. Finally, if $a \notin \overline{K \cap S}$, then $\{a\}$ would be a relatively open and closed subset of the connected set K_S. This proves the asserted closure identity. The intersection of the two closed sectors is L, whose intersection with K is $\{a\}$.

To verify the displayed identity point by point, a point outside K is automatically outside K_S. A point of K outside K_S is in neither S nor L, since $K \cap L = \{a\} \subset K_S$; it therefore lies in S'. Conversely $S' \cap K_S = \varnothing$. If R is chosen with $K \subset B(0, R)$, then

$$S' \setminus \overline{B(0, R)} \subset S' \cap (\mathbb{C} \setminus K), \qquad S' \setminus \overline{B(0, R)} \neq \varnothing.$$

The union-of-connected-sets criterion now has an explicit common point. It does not require $K \cap S'$ to be connected or nonempty; the distinct access condition supplies the latter separately. □

The connectedness part of this proof did not use fullness. If a continuum J meets an embedded line only at a and meets both sides, its closed-side intersections are continua as well. Fullness was used exactly once, in the connected-complement identity. This separates the general-continuum result from the additional hull conclusion.

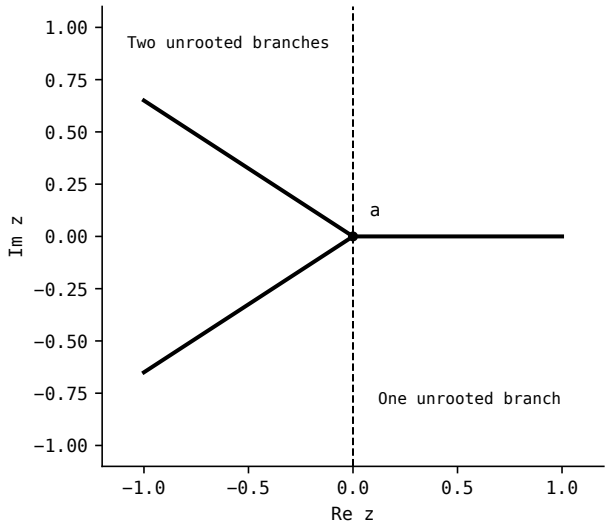


Figure 5.1: A simple cut-line has one root. Each closed-side intersection is a hull; the shared root is retained on both sides. The drawing records this separation convention, not a metric approximation theorem.

5.2 Several disjoint simple cut-lines

Source correspondence. Lyubich, Corollary 9.2(i) and Remark 9.3.

Let $L_1, \dots, L_m$ be pairwise disjoint simple cut-lines with distinct roots $a_1, \dots, a_m$. Pairwise disjointness concerns their plane images; all their spherical closures contain infinity. A cell is a component of $\mathbb{C} \setminus \bigcup_{j=1}^{m} L_j$. There are finitely many cells. A new proper line lies wholly in one old cell and divides that cell into two. This follows either by straightening successively the existing proper lines, or by the Jordan separation theorem applied to the two ends of the new line in that cell. In particular each cell has an unambiguous side relative to each L_j.

Proposition 5.3 (Finite cut cells)**.** *Every cell S meets K, and*

$$K_S = K \cap \overline{S} = \overline{K \cap S} = (K \cap S) \cup \{a_j : a_j \in \partial S\}$$

is a nondegenerate hull. If two different cells are separated by L_j, their subhulls can meet only at a_j.

Proof. Induct on the number of lines. The preceding lemma proves the case of one line. Suppose a new line lies in the old cell S_0. Its root belongs to $K \cap S_0$, since no root lies on a previous line. The old hull $H = K \cap \overline{S_0}$ agrees with K in a neighborhood of this root. The two accesses defining the new cut are therefore distinct for H as well: a connection in arbitrarily small punctured neighborhoods for H would give one for K in the same neighborhoods. Apply Lemma 5.2 to H. Its two closed-side intersections are precisely the new cell intersections. Every other old cell is unchanged. This proves connectedness, fullness and nondegeneracy at the next stage.

The local agreement used in the induction can be expressed without a uniform spacing assumption. If the new root is a, the finite union of old lines is closed and misses a, so there is $r > 0$ with

$$B(a, r) \subset S_0, \qquad H \cap B(a, r) = K \cap B(a, r), \qquad B(a, r) \setminus H = B(a, r) \setminus K.$$

For every $0 < \varepsilon < r$, the two punctured neighborhoods have exactly the same components. Consequently an equality of the two new accesses for H would give equality for K at every

such scale. The contradiction is local at the new root, so previous cuts may be arbitrarily close elsewhere. At each finite stage only finitely many old closed lines must be avoided; the later countable intersection will use compactness instead of demanding a positive distance from their infinite union.

For the closure identity, every point of $K \cap \partial S$ is one of the finitely many boundary roots. Near such a root a_j, all the other lines are absent. In this neighborhood the cell occupies one side of L_j. The one-cut closure assertion therefore supplies points of $K \cap S$ approaching a_j. All other points of K_S already belong to $K \cap S$. The final assertion follows by intersecting the two opposite closed sides of L_j: their common part in K is $K \cap L_j = \{a_j\}$. □

The proposition concerns the subhulls, not the plane closures of the cells. Two plane closures can share an entire cut-line. Keeping the factor K in the intersection prevents a misleading literal interpretation of the source's brief remark about touching closures.

There is also a useful algebraic description. Choose a point $z \in K$ which lies on no cut-line, and let $H_j(z)$ be the closed side of L_j containing z. Then the cell subhull containing z is

$$I_m(z) = K \cap \bigcap_{j=1}^{m} H_j(z).$$

The roots on the boundary are retained by the closed-side convention. A root which lies strictly across some other line is excluded by that other factor. This is why the location of a root on its own line does not make its membership in $I_m(z)$ arbitrary.

5.3 A Jordan cut-curve need not define one cell

Source correspondence. Lyubich, The definition of puzzle pieces in §9.1.1 and Corollary 9.2(ii).

The source also calls a closed Jordan disk bounded by a finite cut-curve a puzzle piece. A finite cut-curve may pass several times through the same hull. That definition alone does not imply that the intersection of the Jordan disk with the hull is connected. The following explicit example records the difference between one finite cut cell and an arbitrary Jordan cut-curve.

Example 5.4 (A disconnected intersection)**.** Set

$$K = [-\tfrac{1}{2}, \tfrac{1}{2}] \cup \{-\tfrac{1}{2} + it : 0 \le t \le 2\} \cup \{\tfrac{1}{2} + it : 0 \le t \le 2\}, \qquad P = \{z : |z - i| \le \tfrac{3}{4}\}.$$

The set K is a simple polygonal arc, so is a full continuum. The circle ∂P meets the two vertical legs at the four points

$$z = \pm\tfrac{1}{2} + i\left(1 \pm \tfrac{\sqrt{5}}{4}\right).$$

All four intersections are transverse and occur in the interiors of straight portions of the arc. Locally the two branches of the circle are on opposite sides of that straight portion. They give distinct accesses, so the circle is a Jordan cut-curve under the source terminology. Nevertheless,

$$K \cap P = \bigcup_{\sigma \in \{-1,1\}} \left\{ \frac{\sigma}{2} + it : 1 - \frac{\sqrt{5}}{4} \le t \le 1 + \frac{\sqrt{5}}{4} \right\}$$

has two components. Its two pieces are a distance 1 apart.

For clarity, substitution on either vertical leg gives the complete inequality rather than only its endpoints:

$$\left| \frac{\sigma}{2} + it - i \right|^2 = \frac{1}{4} + (t-1)^2 \le \frac{9}{16} \quad \Longleftrightarrow \quad |t-1| \le \frac{\sqrt{5}}{4}.$$

The base of the U has imaginary part zero and distance at least 1 from i, so it contributes no connecting segment to $K \cap P$. The two remaining closed intervals have constant real parts $-1/2$ and $1/2$. Any continuous path between them would, by the intermediate value theorem for its real coordinate, contain a point of real part zero. No such point belongs to this intersection. More directly, its intersections with the two open half-planes $\Re z < 0$ and $\Re z > 0$ are a separation into nonempty relatively open sets.

The complement of this polygonal arc is connected, for an arc does not separate the plane; equivalently one can move any point around one of its two free endpoints to the exterior. Thus the example does not exploit a failure of fullness. It refutes the unrestricted connectedness conclusion for this specific source definition. It does not refute Proposition 5.3, whose cells are formed by disjoint *simple* cut-lines and are not arbitrary unions of such cells.

The same configuration tests the unrestricted fjord assertion in source Exercise 9.10(i). The horizontal crosscut $\sigma = [-\frac{1}{2}, \frac{1}{2}] + i$ lies in P and in the complement of K except at its endpoints. Its bounded complementary fjord is

$$V = \{x + iy : -\tfrac{1}{2} < x < \tfrac{1}{2},\ 0 < y < 1\}.$$

For example, $i/10 \in V$ but $|i/10 - i| = 9/10 > 3/4$, so $V \not\subset P$. Thus that containment also needs an additional cell or separation hypothesis; it is not a consequence of an arbitrary Jordan cut-curve.

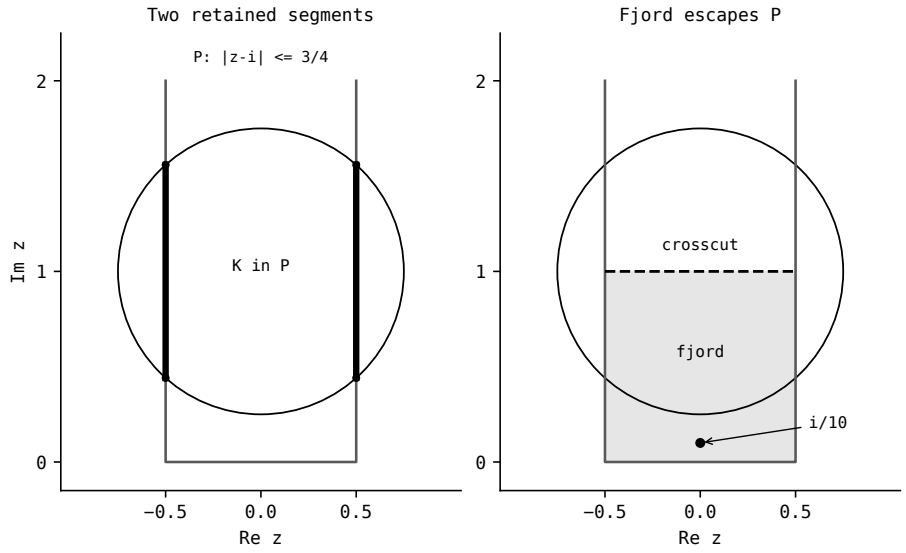


Figure 5.2: Two views of the same U-shaped hull. The disk centered at i with radius $3/4$ meets it in two vertical intervals. The horizontal crosscut at height 1 lies inside the disk, but its bounded rectangular fjord contains $i/10$ outside that disk. Shading or hatching only identifies the regions used in these exact inequalities.

We therefore specify admissibility before deriving a rigidity theorem.

Definition 5.5 (Admissible and perfect pieces)**.** For a continuum K, an admissible Jordan piece is a closed Jordan disk P such that $K \cap P$ is nonempty and connected. It is perfect if $K \cap \operatorname{int} P$ is nonempty and connected. A cut piece has in addition the finite boundary-intersection and access conditions of a Jordan cut-curve. When fullness is needed we require $K \cap P$ to be full as well, or establish that condition separately.

This is a stated normalization, not an additional conclusion extracted from the word "puzzle." For a full K, fullness of $K \cap P$ actually follows automatically: its complement is the union of

$\mathbb{C}\setminus K$ and $\mathbb{C}\setminus P$, two connected sets with nonempty intersection near infinity. Connectedness does not follow from that identity, as the example shows. A compact set can be full and disconnected.

5.4 Solved investigations

Investigation 5.6. For $K=\overline{\mathbb{D}}$, decide whether the tangent line $L=\{1+it : t\in\mathbb{R}\}$ is a cut-line.

Solution. It is a proper embedded line meeting K exactly at 1. However the right half-plane $\{\Re z>1\}$ misses K. Small right semicircles centered at 1 connect the two local branches while avoiding K. For every radius these connections lie in that radius's punctured neighborhood. The accesses are the same, so this is a touching line, not a cut-line. Merely counting $|L\cap K|=1$ does not certify a cut.

Investigation 5.7. Let $K=[-2,2]$ and use the vertical lines $\Re z=-1,0,1$. Determine all cell subhulls and their intersections.

Solution. There are four cells in increasing order of real part. Their subhulls are $[-2,-1]$, $[-1,0]$, $[0,1]$, and $[1,2]$. Each adjacent pair meets at its shared endpoint; every nonadjacent pair is disjoint. The plane closures of adjacent cells share an entire vertical line. Thus the statement about single-point contact is correct only after intersecting with K.

Investigation 5.8. Explain why the complement calculation for $K\cap P$ cannot prove connectedness, even though it proves fullness.

Solution. Take K and P from Example 5.4. Their intersection is the union of two disjoint line segments. Its complement is connected: one may avoid both segments by a path passing around their endpoints. The identity $\mathbb{C}\setminus(K\cap P)=(\mathbb{C}\setminus K)\cup(\mathbb{C}\setminus P)$ therefore gives a true conclusion, but not the conclusion that $K\cap P$ is connected. Connectivity and nonseparation are independent properties of a general compact subset of the plane.

Investigation 5.9. In the proof of Lemma 5.2, why is a relatively clopen set $B\subset K_S\setminus\{a\}$ open in K, and why is the missing root essential?

Solution. Every $b\in B$ lies in the open sector S. Choose an ambient neighborhood V_b with $V_b\cap K_S\subset B$, and replace it by $V_b\cap S$. On this smaller neighborhood, K_S and K agree. Hence $B=K\cap\bigcup_{b\in B}(V_b\cap S)$ is relatively open in K. If B contains a, no neighborhood of a is contained in S; points of K on the opposite side may approach a. The localization step would then fail. Compactness supplies closedness, while the omitted root supplies openness; the two roles are different.

Chapter 6

Countable cutting and the failure of a naive partition

6.1 Retained hulls under infinitely many cuts

Source correspondence. Lyubich, Corollary 9.4.

Let $(L_j)_{j\geq 1}$ be a sequence of pairwise disjoint simple cut-lines for K, with roots a_j. Fix $z \in K \setminus \{a_j : j \geq 1\}$ and use the closed sides $H_j(z)$ from the preceding chapter. The finite retained subhulls form a decreasing sequence:

$$I_m(z) = K \cap \bigcap_{j=1}^{m} H_j(z), \qquad I_{m+1}(z) \subset I_m(z), \qquad z \in I_m(z).$$

There is no assumption of a uniform separation distance between the lines. They may accumulate, and the limiting retained set may have empty interior.

Theorem 6.1 (Countable retained hull)**.** *The set*

$$I(z) = \bigcap_{m\geq 1} I_m(z)$$

is a nonempty full continuum. It is either a nondegenerate subhull of K or a singleton. If no root a_i is separated from z by any other line, then every a_i lies in $I(z)$. In particular $I(z)$ is nondegenerate whenever at least one line is present.

Proof. The finite sets are compact connected sets containing z. Their intersection is nonempty and connected by the nested-continuum theorem.

The compactness argument also gives a useful eventual-containment fact. If G is any open neighborhood of $I(z)$, then

$$I(z) \subset G \quad \Longrightarrow \quad \exists m_0 \; \forall m \geq m_0 : \; I_m(z) \subset G.$$

Otherwise choose $x_m \in I_m(z) \setminus G$ along unbounded indices. A subsequence converges in the compact set K to x. For any fixed r, all sufficiently late selected points belong to the closed set $I_r(z)$, so $x \in I_r(z)$. Hence $x \in I(z)$, while closedness of $K \setminus G$ gives $x \notin G$, a contradiction. If $I(z)$ had a separation into two nonempty compact sets, take disjoint neighborhoods of those sets. Eventual containment would force some connected $I_m(z)$ into their disjoint union while still meeting both, which proves the connectedness assertion too. For completeness, fullness can be checked directly rather than inferred from connectedness. The complements $\mathbb{C} \setminus I_m(z)$ are

connected, increasing with m, and all contain the exterior of one fixed large disk. Consequently their union is connected. De Morgan's law gives

$$\mathbb{C} \setminus I(z) = \bigcup_{m \geq 1} (\mathbb{C} \setminus I_m(z)),$$

so $I(z)$ is full. A root a_i lies on its own boundary line and hence belongs to $H_i(z)$. The nonseparation assumption puts it in every other $H_j(z)$. Thus it belongs to the intersection. As $z \neq a_i$, one retained root already guarantees two distinct points.

Writing the root test as a conjunction prevents an endpoint convention from being lost during the limit:

$$a_i \in I(z) \quad \Longleftrightarrow \quad a_i \in K \ \text{ and } \ \forall j \geq 1,\ a_i \in H_j(z).$$

The factor $j = i$ is automatic because the side is closed. Every other factor is the visible-root hypothesis. In particular the proof does not deduce survival from $a_i \in I_m(z)$ for merely one value of m. For $m < i$ its survival is already a condition imposed by earlier lines, and for $m > i$ it continues to be tested by later ones. Once all factors hold, $\operatorname{diam} I(z) \geq |a_i - z| > 0$ gives an explicit positive lower bound for this particular retained hull, although no bound uniform in the choice of root or cutting system has been asserted. □

The source also gives a description by removed sectors. With S_j the open sector of L_j not containing z, the exact formula is

$$I(z) = K \setminus \bigcup_{j \geq 1} S_j.$$

The factor K is necessary. Replacing it by $\mathbb{C}$ would generally leave an unbounded plane region instead of a subset of the compact hull. This is a set-theoretic normalization of the displayed formula, not a new theorem.

When no root is separated from z by another line, the removed sectors are pairwise disjoint. Indeed, for two disjoint proper lines their opposite-to-z sectors are either disjoint or one contains the other line. In the second case the root of that line would be separated from z. Thus the surviving root condition has a geometric meaning: no removed sector hides another cut behind it.

Proposition 6.2 (Nested admissible pieces)**.** *Let $P_1 \supset P_2 \supset \cdots$ be compact pieces such that $K \cap P_n$ is a nonempty full continuum for every n. Then*

$$I = K \cap \bigcap_{n \geq 1} P_n$$

is a nonempty full continuum. If the pieces all contain a and $\operatorname{diam} P_n \to 0$, then $I = \{a\}$.

Proof. Apply the nested compactness and connectedness argument to $K \cap P_n$. Their complements form an increasing connected union, so the intersection is full. In the shrinking case, any $x \in I$ satisfies $|x - a| \leq \operatorname{diam} P_n$ for every n. Letting $n \to \infty$ gives $x = a$. The point a belongs to all the pieces, so the singleton is actually attained. □

6.2 Accesses which survive visible-root cutting

Source correspondence. Lyubich, Lemma 9.5.

We first record the converse of the access observation used for a simple cut-line. Suppose a proper embedded line L meets a full continuum H only at a, and H meets both open sectors. Then the two branches of L at a represent different accesses to H. Here is a direct justification. Straighten L to the real axis and a to 0. The closed upper and lower intersections of H are

connected by the relative-clopen argument in Lemma 5.2, which did not use the access hypothesis. Fix one nonzero point of H in each half-plane, and choose a small neighborhood of 0 excluding both points. If the two accesses agreed, there would be a path in the complement connecting negative and positive points of the axis inside this neighborhood, chosen to be a convex disk in the straightened coordinates. Approximate the path there by a polygonal path missing H, perturb it to have only finitely many axis contacts, and erase its loops. Among consecutive axis contacts there must be two with opposite signs. The intervening simple excursion lies in one open half-plane. Together with the axis segment between its endpoints it forms a Jordan curve. Its bounded domain lies in the small convex disk, since that disk contains the curve and its convex hull. The excursion is a positive distance from 0, so a sufficiently small half-disk at 0 in the chosen half-plane lies on the bounded side. The closed-side continuum in that half-plane contains 0 and the fixed farther point, but its only possible contact with the Jordan curve is 0. Its intersection with the closed bounded side is therefore relatively open at 0 as well as at its other points, and is relatively closed. It is nonempty and excludes the fixed farther point. This contradicts connectedness of the closed-side continuum.

Proposition 6.3 (Access injection under visible-root removal)**.** *In Theorem 6.1, suppose no root is separated from z by another cut-line. Inclusion $\mathbb{C} \setminus K \subset \mathbb{C} \setminus I(z)$ then induces an injection from accesses to K at z into accesses to $I(z)$ at z.*

Proof. The map is well defined: local connections in the smaller complement are still local connections in the larger one. To prove injectivity, take two different accesses to K at z. By the access and Lindelof results of Volume VIII they are represented by distinct external rays, whose union with z is a simple cut-line M.

We may represent every original cut-line by its Green ray pair without changing which points of K lie on its retained side. Here are the precise winding-number details. Fix an original root a, another point $x \in K \setminus \{a, z\}$, and a small spherical disk about a whose closure misses x and z. The original branch at a and its corresponding Green ray represent the same access. Thus sufficiently short tails can be joined inside this disk by an arc in the complement of K. Truncate both branches at that connector. Their remaining portions, joined through infinity and the connector, form a closed path in $\widehat{\mathbb{C}} \setminus K$. This domain is simply connected, so the path is null-homotopic there. In a chart sending z to infinity its winding number about x is zero. The difference between this truncated path and the full branch comparison is a closed path entirely in the small disk; its winding number about x is also zero. Applying this argument to both branches shows that the original closed cut-curve and its Green representative have identical winding numbers about x. The index of a Jordan curve distinguishes its side containing x from its side containing z, so every retained subset of K is unchanged. The root a itself lies on both curves and is retained. Different roots have disjoint external ray pairs, so the Green representatives remain disjoint. This argument proves side membership for each fixed point; it requires neither a homotopy continuous at the landing root nor a simultaneous ambient isotopy of an infinite family.

One can organize the index computation as a finite algebra of closed paths. Work in the chart sending z to infinity and orient the original and Green cut-curves compatibly. For any closed path C avoiding the fixed point x, use the topological winding number

$$\operatorname{ind}_x C = \deg\left(t \longmapsto \frac{C(t) - x}{|C(t) - x|}\right).$$

This definition requires continuity, not a rectifiable cut-line. Additivity under concatenation and sign change under reversal give

$$\operatorname{ind}_x C_{\text{old}} - \operatorname{ind}_x C_{\text{Green}} = \sum_{i=1}^{2} \epsilon_i(\operatorname{ind}_x A_i + \operatorname{ind}_x E_i) = 0, \qquad \epsilon_i \in \{-1, 1\}.$$

Here A_i is the truncated loop in the spherical exterior, and E_i is the small loop restoring its tail at the root. The former contracts in the exterior; the latter contracts in the chosen disk avoiding x. For a Jordan curve, zero index is exactly the component containing the chart's infinity. Thus the equality compares the retained side relative to z, rather than an arbitrary naming of positive and negative sides.

Let S^+, S^- be the sectors of M. Each removed open sector S_j lies wholly on one side of M. Indeed M is disjoint from its Green boundary, and its root z is not in S_j; the ray-pair separation then prevents S_j from crossing M. If a root a_j lies in S^+, it survives in $I(z)$ by Theorem 6.1, so $I(z) \cap S^+$ is nonempty. If there is no root in S^+, every removed sector lies in S^-, and $K \cap S^+$ is unchanged and nonempty. The same argument applies to the negative side. Thus $I(z)$ meets both sectors of M. The access criterion just proved says that the two branches of M are still different accesses to $I(z)$. This proves injectivity.

For either sign $\tau \in \{+, -\}$ the witness has the explicit alternative

$$\{a_j\} \cap S^\tau \neq \varnothing \implies I(z) \cap S^\tau \supset \{a_j\} \cap S^\tau, \qquad \{a_j\} \cap S^\tau = \varnothing \implies I(z) \cap S^\tau = K \cap S^\tau.$$

In this display $\{a_j\}$ denotes the set of all cutting roots, not one chosen index. The first alternative uses their proved survival. In the second alternative every removed sector is on the opposite side, so no point on this side was removed at all. Either alternative gives a nonempty set. This is why mere nonemptiness of the full intersection $I(z)$ would be inadequate: access distinction requires surviving points on each of two specified sides.

□

The argument proves injection, not surjection. Removing limbs may create new approaches to the retained hull. It also explains the visible-root hypothesis: the surviving root on each affected side is the explicit witness preventing that side from disappearing.

6.3 Nested compactness gives no rate

If P_n are decreasing nonempty compact sets and their intersection is $\{a\}$, their diameters tend to zero. Otherwise one could choose $x_n, y_n \in P_n$ with $|x_n - y_n| \geq \varepsilon$ along a subsequence. Compactness of P_1 supplies limits x, y, and nesting puts both limits in every P_m. They would therefore equal a, contradicting the separation. Thus

$$\bigcap_n P_n = \{a\} \quad \iff \quad \operatorname{diam} P_n \longrightarrow 0 \qquad \text{for a decreasing compact nest containing } a.$$

This equivalence is qualitative. It neither produces a computable depth $n(\varepsilon)$ nor supplies a geometric decay constant. Quantitative rates require additional hypotheses, developed later in the book.

6.4 Why nonseparation need not be transitive

Source correspondence. Lyubich, §9.1.3 and Corollary 9.7.

For a family of cut-lines one can define a symmetric relation by declaring x and y inseparable when no line places them in opposite open sectors. Reflexivity and symmetry are immediate. Transitivity is an additional mathematical property, and it can fail precisely because a point on a cut-line belongs to neither open side.

Example 6.4 (Two tangent disks)**.** Let

$$K = \{z : |z + 1| \leq 1\} \cup \{z : |z - 1| \leq 1\}, \qquad a = 0, \quad x = -1, \quad y = 1.$$

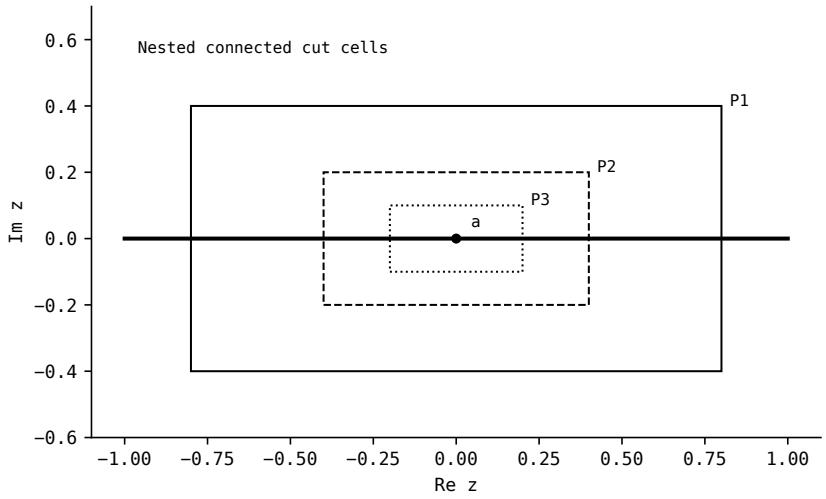


Figure 6.1: A nested family illustrates the distinction between containment and a proved diameter bound. The schematic does not supply a uniform rate for an arbitrary puzzle system.

The hull is locally connected, and the imaginary axis is a simple cut-line through a. It separates x from y. The only possible root of a simple cut-line is a: a line meeting a closed disk in only one point cannot put that connected disk on both sides, and any boundary point other than a has only the exterior access supplied by its circular boundary.

Consequently a cannot be separated from either center. A cut-line through a contains a instead of placing it in a sector; a line avoiding a cannot divide either disk from a while meeting the hull only at a cutpoint. The resulting relation satisfies

$$a \sim x, \qquad a \sim y, \qquad x \not\sim y.$$

It is not an equivalence relation and cannot be the fiber relation of a map.

There is a second way to see the same obstruction directly from Jordan pieces. Any cut-curve for this hull has all of its hull intersections at a. If a Jordan piece contains a in its interior, its boundary avoids K. Connectedness then forces the whole hull to lie inside the piece. Thus the intersection of all such pieces is K. A piece surrounding just the left disk can instead have a on its boundary. These two kinds of intersection overlap without being equal. They cannot be the fibers of one physical-to-puzzle projection.

The source defines an escaping nest by requiring its limiting impression to be either contained in every test piece or disjoint from that piece's interior. In the example a nest whose pieces contain a in their interiors has impression containing K. A piece surrounding one disk violates that escaping condition. The asserted end attached to every physical point is therefore not available from these literal definitions.

There are several mathematically different possible repairs. One may retain an overlap system of impressions; one may restrict the allowable roots and points; or one may construct a closed equivalence relation by an explicitly specified decomposition procedure. These choices need not give the same space. The pinched-disk chapters use a precise closed relation where one has been established, rather than assigning a quotient to a merely symmetric nonseparation relation.

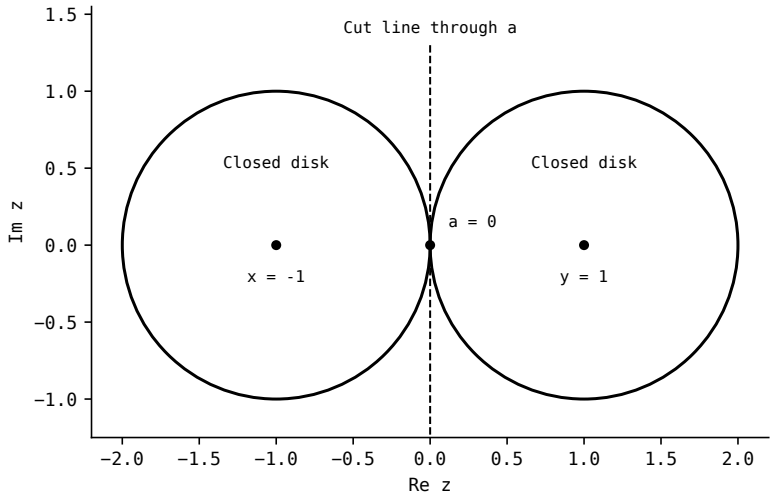


Figure 6.2: Two tangent closed disks expose the endpoint convention: the common root lies on the separating line, whereas the disk centers lie on opposite sides. Inseparability defined with open sides is not transitive.

6.5 A valid conditional projection statement

Proposition 6.5 (Disjoint impressions and quotient semantics)**.** *Let X be compact, let Y be Hausdorff, and let $q : X \to Y$ be a continuous surjection. Suppose a family of nonempty compact sets $(I_y)_{y\in Y}$ satisfies $q^{-1}(y) = I_y$. Then the I_y partition X. If all I_y are singletons, q is a homeomorphism. Conversely, an overlapping family with two unequal members having nonempty intersection cannot be the full fiber family of any map.*

Proof. If $x \in I_y \cap I_{y'}$, then $q(x) = y = y'$, so the two fibers agree. Surjectivity makes every fiber nonempty and their union is X. Singleton fibers make q bijective; a continuous bijection from a compact space to a Hausdorff space is a homeomorphism. For the last assertion, any map has disjoint fibers by the first equality, independent of continuity or topology.

The homeomorphism assertion also has a direct closed-set proof. If $F \subset X$ is closed, compactness of X makes F compact, and continuity makes $q(F)$ compact. Since Y is Hausdorff, $q(F)$ is closed. For bijective q this says

$$(q^{-1})^{-1}(F) = q(F) \ \text{ is closed in } Y \quad \text{for every closed } F \subset X,$$

which is precisely continuity of the inverse. Neither this argument nor the fiber identity constructs a relation from overlapping sets. For the tangent-disk example the obstruction occurs before any topological issue: $q(a) = q(x)$ and $q(a) = q(y)$ would force $q(x) = q(y)$, whereas the proposed relation requires these last values to differ. Passing to a quotient therefore needs an actual repair of the relation, not merely a different topology on the same alleged fibers. □

This elementary result does not construct the map q. It identifies exactly what must be supplied before using a source statement about a projection. The general end construction in the source's earlier topology chapter likewise states a pairwise-disjoint-impression hypothesis. The tangent-disk example shows why that hypothesis cannot be dropped at the puzzle stage.

6.6 Solved investigations

Investigation 6.6. Produce nondegenerate hulls whose decreasing intersection is a point.

Solution. Take $K_n = [-1/n, 1/n]$. Each is a nondegenerate full continuum, $K_{n+1} \subset K_n$, and $\bigcap_n K_n = \{0\}$. Their diameters are $2/n$. Thus a theorem about a nested limit must allow singletons even when no finite stage is a singleton. The example also gives a polynomial rate, not a geometric one.

Investigation 6.7. Given any decreasing positive sequence $r_n \to 0$, construct a compact nest with singleton intersection and diameter exactly r_n.

Solution. The intervals $P_n = [-r_n/2, r_n/2]$ have diameter r_n and intersection $\{0\}$. Choosing $r_n = 1/\log(n+2)$ gives a rate slower than every fixed negative power of n. A singleton intersection alone therefore cannot justify a polynomial or exponential stopping estimate.

Investigation 6.8. For $K = [-2, 2]$, remove the right sector of the vertical line $\Re z = 1$. Compare the formulas with ambient factor K and with ambient factor $\mathbb{C}$.

Solution. With $S = \{\Re z > 1\}$ one has $K \setminus S = [-2, 1]$, a subhull. In contrast, $\mathbb{C} \setminus S = \{\Re z \leq 1\}$ is unbounded and contains points arbitrarily far from K. The error is set-theoretic and cannot be repaired by choosing a different topology on the same formula.

Investigation 6.9. In the tangent-disk example, what does taking the equivalence relation generated by the naive nonseparation relation do to the two centers?

Solution. It identifies them: $x \sim a \sim y$ is a chain of length two. More generally every point of either disk is inseparable from the root by the permitted simple cuts, so the generated equivalence relation collapses the entire hull to one class. This is a possible quotient, but it loses the two-disk geometry and is not the locally connected recognition model. Taking a transitive closure is a substantive choice, not a harmless notation change.

Investigation 6.10. Why is the hypothesis $z \neq a_j$ useful in the last assertion of Theorem 6.1?

Solution. The proof of nondegeneracy uses the distinct retained points z and a_j. If they coincide, that argument gives only one point. Moreover a point on L_j does not choose a unique open side of the line. One must add side data or an independent nondegeneracy argument before allowing the base point to be a root. The hypothesis resolves both the side-selection ambiguity and the final two-point witness.

Chapter 7

Rigidity, local connectivity, and landing

7.1 The topology supplied by a shrinking nest

Source correspondence. Lyubich, Proposition 9.8, Corollary 9.9 and Corollary 9.18.

A set containing a point need not be a neighborhood of that point. To obtain local connectivity from puzzle geometry, the point must lie in the interior of each relevant piece. The connected sets supplied by admissibility then contain relatively open neighborhoods in the continuum.

Definition 7.1 (Admissible rigidity)**.** A point a of a continuum K is admissibly rigid if there are admissible Jordan pieces P_n with

$$a \in \operatorname{int} P_n, \qquad P_{n+1} \subset P_n, \qquad \operatorname{diam} P_n \longrightarrow 0.$$

It is perfectly rigid if the pieces can be chosen perfect. When the pieces are required to be cut pieces, we say cut-rigid or perfectly cut-rigid. These qualifications record which geometric information is available.

Proposition 7.2 (Rigidity implications)**.** *Admissible rigidity implies weak local connectivity at a. Perfect rigidity implies local connectivity at a. If every point of a continuum is admissibly rigid, the continuum is locally connected.*

Proof. Let U be a relatively open neighborhood of a in K. Choose $\varepsilon > 0$ with $K \cap B(a, \varepsilon) \subset U$, and then choose n with $\operatorname{diam} P_n < \varepsilon$. The set $C_n = K \cap P_n$ is connected, contains a, and satisfies

$$K \cap \operatorname{int} P_n \subset C_n \subset U.$$

The left-hand set is an open neighborhood of a in K. Thus C_n is a connected neighborhood, not necessarily an open one. This is weak local connectivity. If the pieces are perfect, $K \cap \operatorname{int} P_n$ itself is connected and open, proving local connectivity at a.

The radii in this argument have different roles. Interior membership supplies some $r_n > 0$ with $B(a, r_n) \subset \operatorname{int} P_n$; the diameter condition supplies the outer bound. Thus

$$K \cap B(a, r_n) \subset K \cap \operatorname{int} P_n \subset C_n \subset K \cap B(a, \varepsilon).$$

There is no requirement that r_n be comparable to ε or to $\operatorname{diam} P_n$. A very narrow piece can still be a neighborhood, provided $r_n > 0$. Conversely, if a were only on ∂P_n, the first inclusion need not hold for any positive radius, even when C_n is connected and tiny. The proof therefore uses both the metric inequality $\operatorname{diam} P_n < \varepsilon$ and the topological assertion of interior membership; neither is a replacement for the other.

For the global assertion, let $V \subset K$ be open and C one of its components. For every $x \in C$, weak local connectivity gives a connected neighborhood N_x of x contained in V. Since N_x meets C, it is contained in C. Its relative interior is an open neighborhood of x inside C. Taking these interiors over $x \in C$ shows that C is open. Components of every open set are therefore open, which is the local connectivity criterion. □

The proof identifies the exact difference between weak and ordinary local connectivity at one point. A connected closed neighborhood gives the first; a connected open neighborhood gives the second. The global implication is not obtained by calling the same neighborhood open. It uses the component criterion at every point.

Proposition 7.3 (An interior body obstructs shrinking cut pieces)**.** *Let K be a compact set and U a nonempty connected open subset of* $\operatorname{int} K$*. Suppose P is a closed Jordan disk with $\partial P \cap K$ finite. If $a \in \overline{U} \cap \operatorname{int} P$, then $U \subset \operatorname{int} P$. In particular* $\operatorname{diam} P \geq \operatorname{diam} U$*.*

Proof. The curve ∂P misses U: any point of U lying on the curve would have a small neighborhood in K, and a nontrivial subarc of the Jordan curve in that neighborhood would give infinitely many intersection points. Since a is an interior point of P and belongs to $\overline{U}$, the set U meets $\operatorname{int} P$. Connectedness of U, together with $U \cap \partial P = \varnothing$, puts all of U in that component of the curve complement. The diameter inequality follows by taking suprema over pairs of points in U. □

For hull-cut pieces this excludes rigidity at every point of the closure of a nondegenerate interior component. It does not exclude local connectivity there. A closed disk is locally connected at every point, but its boundary cannot be crossed by a finite hull-cut curve in order to make arbitrarily small hull-cut pieces. Dipole cuts of the boundary continuum are a different construction and will be treated separately.

Example 7.4 (A nest which retains a comb impression)**.** Consider the full continuum

$$K = ([0,1] \times \{0\}) \cup (\{0\} \times [0,1]) \cup \bigcup_{n \geq 1} (\{1/n\} \times [0,1]).$$

The rectangles $P_n = [-1/n, 1/n] \times [-1/n, 1 + 1/n]$ are nested, their intersections with K are connected, and their limiting intersection with K is the entire left segment $\{0\} \times [0,1]$. At $a = (0, 1/2)$ the hull is not even weakly locally connected. Any neighborhood of a contains points on infinitely many teeth. A connected subset of the ball $B(a, 1/4)$ cannot join a tooth to the left segment, since all such connections in K must pass through the base at height 0. Hence no connected neighborhood of a lies in that ball. The nonshrinking impression records a real obstruction; it is not a failure to choose a fast enough depth.

7.2 Finite barriers in the prime-end disk

Source correspondence. Lyubich, Exercise 9.10 and Corollary 9.11.

Let K be a hull and let $\Phi : \{w : |w| > 1\} \to \mathbb{C} \setminus K$ be its normalized exterior uniformization. Its prime-end boundary is the unit circle. For $\zeta \in \mathbb{T}$ write $I(\zeta)$ for the cluster set of Φ at ζ from the exterior. The pole at infinity has no role in the local boundary argument below.

Lemma 7.5 (A finite cut-curve away from its endpoint angles)**.** *Suppose P is a bounded Jordan disk whose boundary meets K in finitely many points and whose complementary arcs lie in $\mathbb{C} \setminus K$. Let $B(P) \subset \mathbb{T}$ be the finite set of prime-end angles associated with the accesses of these arcs at their endpoints. If $\zeta \notin B(P)$ and $I(\zeta)$ meets* $\operatorname{int} P$*, then $I(\zeta) \subset P$.*

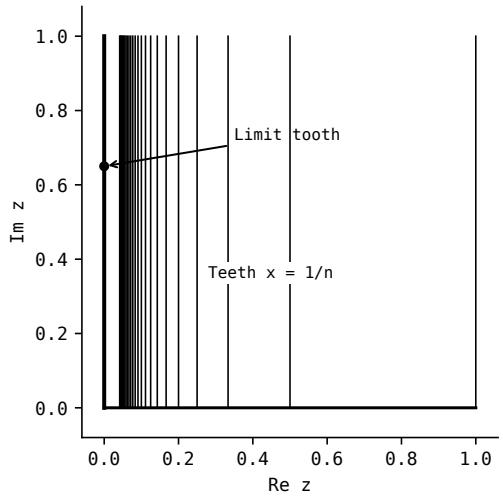


Figure 7.1: A finite prefix illustrates the infinite comb in Example 7.4. The proof uses the whole sequence of teeth and its limit segment; a finite drawing alone cannot establish non-local-connectivity.

Proof. Each open arc of $\partial P \setminus K$ pulls back under Φ^{-1} to an arc in the exterior disk. The inverse-map landing theorem from Volume VIII gives a limit angle at each finite endpoint. Thus the union of these pulled-back arcs and their endpoint angles is a compact subset of the closed exterior disk, apart from no point at infinity because ∂P is bounded. Its intersection with $\mathbb{T}$ is exactly $B(P)$.

Denote that compact pulled-back set by A_P. Since it misses the chosen point, its distance gives an explicit available scale:

$$d = \operatorname{dist}(\zeta, A_P) > 0, \qquad 0 < \delta < \min\{d/2, 1/2\}.$$

The compactness here is essential: disjointness from a nonclosed union of arcs would not by itself give positive distance. There are only finitely many boundary arcs, and adjoining each of their two landing angles makes each image compact. Boundedness of ∂P and properness of the exterior coordinate near infinity prevent an additional escape to infinity. With this choice the same neighborhood avoids every boundary arc at once.

If $\zeta \notin B(P)$, choose $\delta > 0$ so small that

$$V_\delta = \{w : |w| > 1,\ |w - \zeta| < \delta\}$$

is disjoint from the pulled-back boundary. It is connected. Hence $\Phi(V_\delta)$ is connected and avoids ∂P, so it lies wholly in one complementary component of that curve. A sequence witnessing a point of $I(\zeta) \cap \operatorname{int} P$ eventually enters $\operatorname{int} P$. Therefore the entire image lies inside P. Any cluster point as $w \to \zeta$ then belongs to the closed disk P.

Equivalently, if $\eta_k \downarrow 0$ and $\eta_k < \delta$, the cluster definition gives

$$I(\zeta) = \bigcap_k \overline{\Phi(V_{\eta_k})} \subset \overline{\Phi(V_\delta)} \subset P.$$

The last inclusion uses that P is closed. It includes sequences approaching tangentially or wandering inside the cap, not just points on the radial ray. A radial landing statement alone would control a single approach and would not imply this inclusion for the full impression. The sequence witnessing an interior point decides the side of the boundary curve; connectedness then controls all other approaches inside the same cap.

To check connectedness directly, rotate so that $\zeta = 1$ and write $w = 1 + \rho e^{it}$, where $0 < \rho < \delta$. The exterior condition is

$$|1 + \rho e^{it}|^2 > 1 \quad \Longleftrightarrow \quad 2\cos t + \rho > 0.$$

At fixed ρ, the allowed angular interval containing $t = 0$ is $(-\arccos(-\rho/2), \arccos(-\rho/2))$. Moving the argument toward zero stays in this interval, and then moving along the positive radial segment connects to any fixed point $1 + \rho_0$ with $0 < \rho_0 < \delta$. Thus every point can be joined to one reference point inside V_δ. Rotating back proves path connectivity without a convexity assumption on the exterior cap.

□

The endpoint exclusion is essential to this proof. At an angle of a boundary arc, one cannot choose a full angular neighborhood avoiding that arc. The source's statements about impressions touching a closed puzzle piece require additional endpoint analysis; the lemma does not erase that obligation. Shrinking pieces around an interior point provide a useful way to handle the exceptional angles explicitly.

Theorem 7.6 (Landing at a cut-rigid point)**.** *Let $a \in K$ admit shrinking cut pieces P_n with $a \in \operatorname{int} P_n$. If $a \in I(\zeta)$, then $I(\zeta) = \{a\}$. In particular the external ray of angle ζ lands at a. There is at least one such angle.*

Proof. The pieces need not be assumed admissible for this particular conclusion: the finite barrier and shrinking assumptions are the ingredients used. For a fixed ζ, suppose it belongs to $B(P_n)$ for infinitely many n. At each such index the corresponding boundary arc lands at a vertex $b_n \in \partial P_n \cap K$. Lindelof's theorem says that the external ray of angle ζ lands at b_n. A ray has at most one limit, so all these b_n are a single point b. Since $a \in P_n$ and $\operatorname{diam} P_n \to 0$,

$$|b - a| \le \operatorname{diam} P_n \longrightarrow 0$$

along those indices, giving $b = a$. But a is in the interior of every P_n, and cannot be its boundary vertex. This contradiction proves that $\zeta \notin B(P_n)$ for all sufficiently large n.

In quantified form, the conclusion is

$$E_\zeta = \{n : \zeta \in B(P_n)\} \text{ is finite}, \qquad \exists N_\zeta \ \forall n \ge N_\zeta : \ \zeta \notin B(P_n).$$

The index N_ζ may depend on the angle. No uniform bound on the number of exceptional indices is being extracted from finiteness of each $B(P_n)$. Infinitely many finite sets can contain a fixed point infinitely often; the decisive extra information is that every occurrence forces the same radial landing value. Along unbounded indices the diameter estimate forces that value to be a, whereas the strict interior condition forbids a from being any of the boundary vertices. This establishes the eventual exclusion without an assumed rate of shrinking.

For those indices the cluster set meets $\operatorname{int} P_n$ at a, so Lemma 7.5 gives $I(\zeta) \subset P_n$. Intersecting and using the shrinking diameters yields $I(\zeta) \subset \{a\}$. The assumed membership gives equality. A singleton cluster set is equivalent to a continuous boundary limit, and therefore includes the radial limit.

To prove existence, first note that a cannot be in $\operatorname{int} K$: Proposition 7.3 applied to a small disk around a would prevent shrinking. Choose $z_j \in \mathbb{C} \setminus K$ tending to a. The preimages $w_j = \Phi^{-1}(z_j)$ stay bounded and have a subsequence tending to the unit circle. They cannot have a limit of modulus greater than one, since continuity of Φ there would put a in the exterior. They cannot tend to infinity, by the exterior normalization. A subsequence therefore tends to some $\zeta \in \mathbb{T}$ with $a \in I(\zeta)$. □

This proof uses neither a global puzzle quotient nor a presumed partition of all physical points. Its essential finite information is the endpoint set $B(P_n)$. The eventual exclusion follows from uniqueness of a radial limit and the strict interior condition, which is why replacing $a \in \operatorname{int} P_n$ by $a \in P_n$ would invalidate the argument.

Corollary 7.7 (The complete landing fiber)**.** *For a cut-rigid point a, the set*

$$Q(a) = \{\zeta \in \mathbb{T} : I(\zeta) = \{a\}\}$$

is nonempty and compact. The map Φ extends continuously to $\{|w| > 1\} \cup Q(a)$ with value a on $Q(a)$, and no other prime-end impression contains a.

Proof. Existence and the last assertion follow from the theorem. If $\zeta_j \in Q(a)$ tends to ζ, choose w_j outside the disk with $|w_j - \zeta_j| < 1/j$ and $|\Phi(w_j) - a| < 1/j$. Then $a \in I(\zeta)$, so the theorem puts ζ in $Q(a)$. This proves closedness and hence compactness. At each point of $Q(a)$ the singleton-cluster criterion gives continuity from the exterior, and sequences within $Q(a)$ have constant image a. These are exactly the sequences needed for continuity on the stated union.

Compactness of $Q(a)$ even gives a uniform qualitative boundary estimate:

$$\forall \varepsilon > 0\ \exists \delta > 0 : \quad |w| > 1,\ \operatorname{dist}(w, Q(a)) < \delta \implies |\Phi(w) - a| < \varepsilon.$$

If this failed, choose w_j with distance to $Q(a)$ tending to zero but $|\Phi(w_j) - a| \geq \varepsilon$. Choose nearest points $\zeta_j \in Q(a)$ and pass to a subsequence with $\zeta_j \to \zeta \in Q(a)$. Then $w_j \to \zeta$, contradicting the singleton cluster set at ζ. This estimate expresses continuity on a compact set of landing angles. It supplies neither an explicit modulus $\delta(\varepsilon)$ nor a global boundary extension at angles outside $Q(a)$, so it does not enlarge the theorem to the unresolved general quotient setting. □

7.3 A local-connectivity criterion using finite sectors

Proposition 7.8 (Connected open cut neighborhoods)**.** *Let $a \in K$. Suppose that for every $\varepsilon > 0$ there are finitely many cut-lines avoiding a such that the cell S containing a satisfies*

$$K \cap S \text{ is connected}, \qquad K \cap S \subset B(a, \varepsilon).$$

Then K is locally connected at a.

Proof. The cell S is open, so $K \cap S$ is a relatively open neighborhood of a. The assumed connectedness and radius bound provide a basis of connected open neighborhoods. This is precisely local connectivity. □

The radius bound can often be obtained by a finite-cover argument. If each $x \in K \setminus B(a, \varepsilon)$ lies in an open sector separated from a, those sectors cover that compact set, so finitely many suffice. What remains to be proved in a particular application is that the selected lines form a compatible cell system and that removing the boundary roots leaves its interior intersection connected. The terminology "well branched" is designed to supply that last connectivity in suitable systems. It is not a replacement for checking compatibility of the cuts. Similarly, constructing a small Jordan puzzle from the resulting cell is an additional geometric step if perfect *cut-rigidity*, rather than local connectivity, is the target conclusion.

7.4 Solved investigations

Investigation 7.9. Why is the condition $a \in \operatorname{int} P_n$ indispensable in the rigidity definition?

Solution. Take $K = [-1, 1]$, $a = 0$, and P_n small closed disks tangent to K at 0 from above. Their diameters tend to zero and $K \cap P_n = \{0\}$, but $K \cap \operatorname{int} P_n$ is empty. The sets $K \cap P_n$ are not neighborhoods of a in K. Thus this construction gives no neighborhood basis. Membership in a shrinking compact set is weaker than membership in its interior, even when the intersection is a singleton.

Investigation 7.10. Can a point of the unit circle be cut-rigid for the closed unit disk under the finite hull-cut convention?

Solution. No. Apply Proposition 7.3 with $U = \mathbb{D}$. Every piece containing that point in its interior and meeting the hull only finitely on its boundary contains $\mathbb{D}$, so has diameter at least 2. This is consistent with the disk being locally connected. Cut-rigidity is a sufficient geometric criterion in its stated scope, not a necessary condition for local connectivity of arbitrary full hulls with interior.

Investigation 7.11. In Theorem 7.6, may the sets $B(P_n)$ be disjointly replaced by one uniformly bounded exceptional set?

Solution. No such uniform bound is assumed: the number and angles of vertices may vary with n. The proof instead fixes one angle ζ and shows that it cannot occur in infinitely many $B(P_n)$. This pointwise eventual exclusion suffices for the fixed cluster set $I(\zeta)$. It does not give a depth uniform over every angle, and the theorem claims none.

Investigation 7.12. Explain why compactness of $Q(a)$ is not obtained merely by saying that radial limits depend continuously on the angle.

Solution. For a general hull radial limits need not vary continuously in angle, and may not exist at every angle. The proof uses a diagonal sequence $w_j \to \zeta$ with $\Phi(w_j) \to a$, which puts a in a cluster set. Cut-rigidity then forces that entire cluster set to be $\{a\}$. The continuity conclusion is obtained after this argument, not assumed before it.

Investigation 7.13. State the exact compact set used when replacing pointwise separating sectors by finitely many sectors in Proposition 7.8.

Solution. Use $K \setminus B(a, \varepsilon)$, which is closed in the compact set K. Each of its points must belong to an *open* removed sector. Compactness gives a finite subcover. The complement in K of that union then lies in $B(a, \varepsilon)$. A cover by closed sectors would not give the same finite-subcover conclusion without an additional argument, and using $K \setminus \overline{B(a, \varepsilon)}$ would lose compactness.

Chapter 8

Gates and Jordan interior components

8.1 A gate onto one interior component

Source correspondence. Lyubich, Proposition 9.21 and its dependency on source Lemma 1.30.

Let K be a locally connected hull and U a component of $\operatorname{int} K$. The uniform small-arc property of a compact locally connected continuum is available from the foundational chapters. In particular any two points can be joined by an arc in K, and nearby points can be joined by an arc of arbitrarily small diameter. No assertion about shrinking interior components is needed for the construction that follows.

First observe that

$$\partial U \subset \partial K.$$

Indeed, if a boundary point of U had a small disk contained in $\operatorname{int} K$, that connected disk would meet U and hence lie in the same component U. This would make the point an interior point of U, a contradiction.

Theorem 8.1 (Gate retraction)**.** *For $x \in K \setminus \overline{U}$, choose an arc in K from x to $\overline{U}$, and let $g_U(x)$ be its first intersection with $\overline{U}$. This point is independent of the chosen arc. Setting $g_U(x) = x$ on $\overline{U}$ defines a continuous retraction*

$$g_U : K \longrightarrow \overline{U}$$

which is constant on every component of $K \setminus \overline{U}$. If an arc $\alpha \subset K$ meets $\overline{U}$, then $g_U(\alpha) \subset \alpha$. If it does not meet $\overline{U}$, its image under g_U is a singleton.

Proof. Suppose two arcs from x have different first contacts u_1, u_2. Truncate them at those contacts. Their interiors avoid $\overline{U}$. Take the final point v of the first truncated arc which belongs to the second arc. Such a point exists by compactness of the intersection. The subarcs from v to u_1 and u_2 have disjoint interiors and meet only at v, which lies outside $\overline{U}$.

More explicitly, parametrize the first truncated arc injectively by $\alpha_1 : [0,1] \to K$, with endpoint u_1, and write A_2 for the image of the other. The set of contact parameters is compact and nonempty, so

$$t_* = \max\{t \in [0,1] : \alpha_1(t) \in A_2\}, \qquad v = \alpha_1(t_*).$$

Since $u_1 \neq u_2$ and the second arc has no other contact with $\overline{U}$, one has $t_* < 1$. Every point after t_* on the first arc misses the entire second arc. Taking the subarc of A_2 from v to u_2 therefore produces exactly the required fork. This selection uses order along an embedded arc

and compactness of its intersection, rather than an unproved assumption that two arbitrary paths have a last common geometric segment.

Choose points $w_i \in U$ close to u_i. By the small-arc property, join u_i to w_i by arcs $\eta_i \subset K$ of sufficiently small diameter. The compact branch $[v, u_2]$ misses u_1, and the compact branch $[v, u_1]$ misses u_2. The two contacts are distinct. We may therefore choose the arcs small enough that η_1 misses the other branch and η_2, and η_2 misses the first branch and η_1. Erasing any loops in $[v, u_i] \cup \eta_i$ gives an arc δ_i from v to w_i; the two arcs meet only at v.

The smallness choices can be made simultaneously. The two positive distances from a contact to the opposite compact branch, together with $|u_1 - u_2|$, have a positive minimum d. Choose the connecting arcs with diameter less than $d/4$. Then

$$\eta_i \subset B(u_i, d/4), \qquad B(u_1, d/4) \cap B(u_2, d/4) = \varnothing, \qquad \eta_i \cap [v, u_{3-i}] = \varnothing.$$

To obtain such an arc, first use the uniform small-arc property to choose the permissible distance from w_i to u_i, and only then choose $w_i \in U$ within that distance. Boundary membership $u_i \in \partial U$ guarantees this last choice. The order of these choices ensures that the added arcs do not accidentally meet the opposite branch while joining the contacts to the open component.

Join w_1 to w_2 by a path in the open connected domain U, and erase loops to make it an arc. Trim this connector at its last contact with δ_1 before its first subsequent contact with δ_2. The trimmed arc lies in U, has one endpoint on each δ_i, and its interior misses both. The two initial branch segments from v to those endpoints, together with the connector, form a Jordan curve $\Gamma \subset K$. Its bounded complementary domain V is contained in K by fullness. Since V is open, it lies in $\operatorname{int} K$. Moreover it meets U: near any interior point of the connector, a small disk lies in U and meets the bounded side of Γ. Connectedness therefore gives $V \subset U$. But $v \in \partial V$, so $v \in \overline{U}$, contradicting its construction. The first contact is independent of the arc.

The component argument at the contradiction is worth separating from the filling argument. Fullness first gives $V \subset K$; openness then gives $V \subset \operatorname{int} K$. The set V is connected and has a point in U, so maximality of a connected component implies

$$V \cap U \neq \varnothing, \quad V \subset \operatorname{int} K \quad \Longrightarrow \quad V \subset U \quad \Longrightarrow \quad \Gamma = \partial V \subset \overline{U}.$$

The final inclusion contradicts the particular vertex $v \notin \overline{U}$. It is not enough to observe that the two tentative contacts could be joined inside U: that observation alone would not compare paths outside U. The Jordan curve forces its entire bounded region into a single interior component, which is the step excluding two distinct gates.

The open subspace $K \setminus \overline{U}$ is locally path connected. Every one of its components is therefore path connected. If x and y lie in one such component, join them by a path avoiding $\overline{U}$ and append an arc from y to its gate. Loop erasure gives an arc from x whose only contact with $\overline{U}$ is that same endpoint. Thus $g_U(x) = g_U(y)$. Components of the open subspace are open in K, so this also proves local constancy and continuity away from $\overline{U}$.

At a point $a \in \overline{U}$, let $x \in K$ be close to a. A small arc α joining x to a has a first contact with $\overline{U}$ if x lies outside it. That contact is $g_U(x)$ and lies on α. If x lies in $\overline{U}$, its image is x. In either case

$$|g_U(x) - a| \leq \operatorname{diam} \alpha.$$

The small-arc property proves continuity at a.

Indeed, given $\varepsilon > 0$, choose $\delta > 0$ from the small-arc property of K so that points at distance less than δ admit an arc of diameter less than ε. The gate just identified belongs to that arc, so

$$x \in K, \quad |x - a| < \delta \quad \Longrightarrow \quad |g_U(x) - g_U(a)| = |g_U(x) - a| < \varepsilon.$$

This handles approaches through arbitrarily many different components of $K \setminus \overline{U}$. Local constancy alone would only handle a sequence eventually remaining in one fixed component. The small-arc

estimate supplies the missing control as these components accumulate at the retraction's fixed set.

Finally let an arc α meet $\overline{U}$. A point of the arc outside $\overline{U}$ belongs to an open parameter interval whose endpoint, in at least one direction, is a contact with $\overline{U}$. The subarc to that contact gives its gate, which lies on α. Points already in $\overline{U}$ are fixed. Hence $g_U(\alpha) \subset \alpha$. If the arc misses $\overline{U}$, it lies in one component of the complement and the already proved constancy applies. □

The first-contact assertion is stronger than choosing an arbitrary path from the hull to the component. It makes the retraction canonical once U is specified. The proof's only uniform input was the small-arc property of K. In particular it does not depend on the null-family theorem that will use these gates.

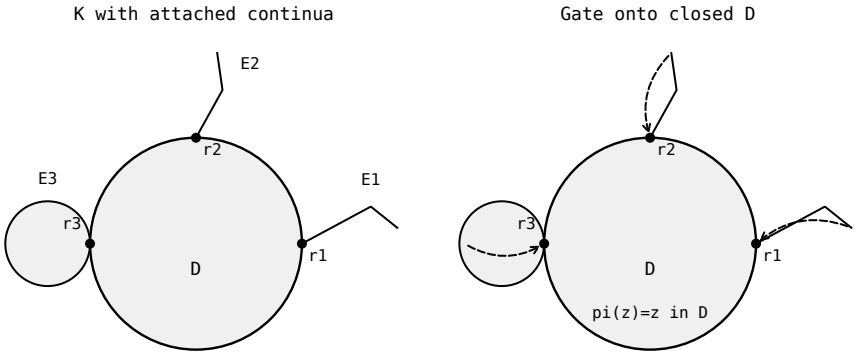


Figure 8.1: A disk body with attached continua illustrates a gate retraction. The body is fixed pointwise; each indicated attachment maps to its marked root. Fibers need not have small diameter. The theorem establishes continuity from small connecting arcs in the whole hull.

Corollary 8.2 (At most one common boundary point)**.** *If U and V are distinct components of* $\operatorname{int} K$*, then $\overline{U} \cap \overline{V}$ is empty or a singleton.*

Proof. The open set V misses $\overline{U}$. Otherwise a disk in V around a point of $\overline{U}$ would meet U, joining two supposedly different components of $\operatorname{int} K$. The connected set V therefore lies in one component of $K \setminus \overline{U}$, so g_U is constant on it. Continuity makes g_U constant on $\overline{V}$. Every point of $\overline{U} \cap \overline{V}$ is fixed by the retraction. Hence there can be at most one such point.

If the constant value on V is c, the two properties of the retraction give the useful pair of equations

$$g_U|_{\overline{V}} \equiv c, \qquad g_U|_{\overline{U}} = \mathrm{id}_{\overline{U}}\,.$$

For any $p, q \in \overline{U} \cap \overline{V}$, substitution yields $p = g_U(p) = c = g_U(q) = q$. This is a statement about closures, although constancy was initially proved only on the open component V. Continuity is exactly what permits passage to its boundary. Disjointness of the two open components by itself would not limit their common boundary to one point; many planar domains with disjoint interiors have long common boundary arcs. □

This proves the conclusion of source Exercise 9.22 before proving that the components are Jordan disks. It is not circular: the gate theorem used arcs in K, not internal ray coordinates or a pinched-disk model.

8.2 The boundary of an interior component

The exterior Caratheodory theorem gives a continuous surjection $\Phi : \mathbb{T} \to \partial K$ because K is locally connected. Composing with the gate gives

$$\mathbb{T} \xrightarrow{\Phi} \partial K \xrightarrow{g_U} \partial U.$$

The second arrow has precisely the indicated range. A point of $\partial K \cap \overline{U}$ lies in ∂U, and a point outside $\overline{U}$ is sent to a first-contact point in ∂U. Conversely every point of ∂U belongs to ∂K and is fixed. Thus ∂U is a continuous image of the circle and is locally connected. This argument avoids the invalid general inference that local connectivity of the closure of a domain alone implies local connectivity of its boundary.

Theorem 8.3 (Jordan interior components)**.** *Every component U of the interior of a locally connected hull K is a Jordan domain. Its Riemann map $\psi : \mathbb{D} \to U$ extends to a homeomorphism $\overline{\mathbb{D}} \to \overline{U}$.*

Proof. First U is simply connected. A Jordan curve contained in U bounds a disk contained in K by fullness. That open disk meets U near the curve, so belongs to the same component of $\operatorname{int} K$ and is contained in U. The planar simple-connectivity criterion applies. Since U is bounded and nonempty, it has a Riemann map $\psi : \mathbb{D} \to U$. The preceding gate argument proves local connectivity of ∂U, so Caratheodory's theorem extends ψ continuously to the closed disk.

Suppose two distinct circle points ξ_1, ξ_2 have the same image $a \in \partial U$. Their two radial arcs have disjoint images except for the common center $b = \psi(0)$ and the common landing point a. Their union is a Jordan curve $\Gamma \subset \overline{U} \subset K$. Let V be its bounded complementary domain. Fullness gives $V \subset K$, and V meets U along the interior side near a point on either radial arc. Thus $V \subset U$.

This side choice is a separation statement inside the open domain U. Removing the two radial images and their common center corresponds under ψ to removing the two radii from $\mathbb{D}$. The two remaining angular sectors are connected. Each image avoids Γ, so is contained in one side of Γ; near an interior point of a radial image, the local homeomorphism ψ sends the two sides of the radius to the two sides of the Jordan arc. Thus the two sector images occupy opposite sides, and exactly one lies in the bounded side V:

$$\psi(W_1) \subset V, \qquad \psi(W_2) \subset \mathbb{C} \setminus \overline{V}$$

after labeling the sectors appropriately. This justifies selecting the circle interval adjacent to W_1 for the following boundary-limit argument.

One of the two angular sectors between the radii maps into V. Its boundary circle interval A has nonempty interior in $\mathbb{T}$. Continuity gives $\psi(A) \subset \overline{V} \cap \partial U$. The interior of V is contained in U, and $\Gamma \cap \partial U = \{a\}$, so ψ is identically a on A. This is impossible for a nonconstant holomorphic function. Indeed, on a smaller open circle arc the function $\psi - a$ has continuous real boundary value zero; Schwarz reflection across that circle arc extends it holomorphically, and the identity theorem then makes it identically zero. Equivalently the reflection assertion follows from Morera's theorem by splitting small triangles across the arc, as in the boundary-uniqueness argument of Volume VIII.

For an explicit local formulation, take a conformal coordinate τ from a small upper half-disk to the disk side of the selected circle arc, and put $h = \psi \circ \tau - a$. It is holomorphic above the real diameter, continuous on it, and zero there. The reflection is

$$\widetilde{h}(w) = \begin{cases} h(w), & \Im w \geq 0, \\ \overline{h(\overline{w})}, & \Im w < 0. \end{cases}$$

Its two formulas agree on the real diameter. Schwarz reflection, or the triangle-splitting Morera proof cited above, makes it holomorphic across that diameter. Since its zeros include an interval with an interior accumulation point in the enlarged domain, the identity theorem gives $\widetilde{h} \equiv 0$. Hence ψ is constant on an open subset of $\mathbb{D}$ and then on $\mathbb{D}$, contradicting conformality. No differentiability of the original boundary parametrization is required.

The extension is therefore injective on $\mathbb{T}$. Interior and boundary images cannot coincide, and the original Riemann map is injective inside. It is a continuous bijection from the compact closed disk to the Hausdorff space $\overline{U}$, hence a homeomorphism. In particular ∂U is a Jordan curve. □

A bounded simply connected domain with a continuous boundary extension need not have an injective extension. The extra fullness of the ambient hull was used to force the Jordan region V into the same interior component. That is the step which rules out slit-type identifications.

8.3 The null-family consequence

The foundational null-family theorem, whose geometric input is Theorem 8.1, now applies to the interior components:

$$\forall \varepsilon > 0 : \quad \#\{U \text{ a component of } \operatorname{int} K : \operatorname{diam} U \geq \varepsilon\} < \infty.$$

It is useful to record why the gate's arc-image assertion gives a uniform estimate. If nearby points x, y can be joined by an arc of diameter less than ε, then either the gate is constant on that arc or both gate images lie on it. Thus the same modulus works for every component:

$$|x - y| < \delta(\varepsilon) \quad \Longrightarrow \quad |g_U(x) - g_U(y)| < \varepsilon \qquad \text{for every } U.$$

This uniformity, not a separate modulus chosen for each component, is what prevents infinitely many large interior components. The complete compactness argument is canonical in Theorem 3.5; no second null-family claim is being inferred merely from disjointness or finite area. Disjoint thin domains can have large diameters and arbitrarily small areas.

The same uniformity controls compact images, which is useful when taking limits of components or arcs. For nonempty compact $E, F \subset K$, the gate modulus implies

$$d_H(E, F) < \delta(\varepsilon) \quad \Longrightarrow \quad d_H(g_U(E), g_U(F)) \leq \varepsilon \qquad \text{for every } U.$$

To check one directed distance, choose $x \in E$ and $y \in F$ with $|x-y| < \delta(\varepsilon)$. Then $|g_U(x)-g_U(y)| < \varepsilon$. Take the supremum over x, and interchange E, F for the reverse distance. The same δ works for all U, so a component index may vary along a convergent sequence. This is a consequence of the canonical gate estimate; it is not a separate geometric shrinking assumption.

The gates also clarify attachment points. A connected component of $K \setminus \overline{U}$ has one gate, and its closure can meet $\overline{U}$ only at that gate. Indeed continuity makes its gate value constant on its closure, while points of $\overline{U}$ are fixed. Distinct attachment continua may share a gate; the statement does not assert a bound on their number or their diameters without further input.

8.4 Solved investigations

Investigation 8.4. For two tangent closed disks, describe the gate onto the left disk.

Solution. It fixes every point of the left disk and sends every point of the right disk outside the tangency point to that tangency point. Any arc from the right disk to the left must pass through the common point. The resulting map is continuous because a sequence in the right

disk can approach a point of the left disk only at the tangency point. This simple retraction is compatible with the failure of the naive nonseparation partition: a retraction onto one selected component is not a simultaneous partition by all possible puzzle impressions.

Investigation 8.5. Suppose every pair of points at distance less than δ can be joined in K by an arc of diameter at most 3ε. Bound their gate images for every interior component.

Solution. The gate image of the joining arc is either a singleton or a subset of the arc. Consequently $|g_U(x) - g_U(y)| \leq 3\varepsilon$, uniformly in U. No length estimate is needed: the joining arc may have infinite length. The relevant control is its diameter. This distinction permits the argument on locally connected fractal hulls without rectifiability.

Investigation 8.6. Can two distinct interior components of a locally connected hull share a nontrivial boundary arc?

Solution. No. The gate onto the closure of the first component is constant on the second component and, by continuity, on its closure. It fixes their common boundary points. A shared arc would contain at least two such fixed points, contradicting constancy. This proof also rules out two isolated common boundary points and does not require the common set itself to be connected.

Investigation 8.7. Give disjoint bounded plane domains of area tending to zero whose diameters do not tend to zero. Explain why this does not contradict the null-family theorem for a locally connected hull.

Solution. Take disjoint horizontal rectangles of width 1 and heights tending to zero, stacked inside a bounded square with disjoint vertical intervals. Every diameter is at least 1, while their areas tend to zero. Their closures together with their limit segment do not supply the required uniform small-arc geometry across the intervening gaps. A suitable filled locally connected ambient hull would join those regions through its interior and change the component list. The theorem concerns components of one locally connected full hull, not an arbitrary disjoint family.

Chapter 9

Dipoles, attachments, and peripheral geometry

9.1 Dipoles cut a boundary differently from a hull

Source correspondence. Lyubich, §9.2.1, Exercises 9.16–9.17 and Lemma 9.19.

Let J be a planar continuum which need not be full. A dipole cut-line meets J at two points, and the segment between them lies in a bounded component of $\mathbb{C} \setminus J$. For a Jordan curve this is the familiar construction of an interior crosscut continued outside the curve at both ends. It is a cut of the boundary continuum, not a simple cut of its filled hull: the middle segment lies inside that hull.

A single dipole line for a Jordan curve divides the curve into two arcs. It is tempting to repeat the finite simple-cut induction without checking what the middle cells contain. That induction is not valid for arbitrary families of dipoles.

Example 9.1 (Two dipoles and a disconnected middle cell)**.** Take $J = \mathbb{T}$ and the two lines

$$L_- = \{\Im z = -\tfrac{1}{2}\}, \qquad L_+ = \{\Im z = \tfrac{1}{2}\}.$$

Each line meets J transversely in two points, and the segment between those points is contained in $\mathbb{D}$. They are disjoint dipole lines. For the middle cell $S = \{z : -1/2 < \Im z < 1/2\}$,

$$J \cap \overline{S} = \{e^{it} : |\sin t| \le \tfrac{1}{2}\}$$

is the union of two disjoint circular arcs, one containing 1 and the other containing -1. Thus the general dipole version of the connected cell conclusion in source Exercise 9.16 requires a restriction.

The source sentence uses both “simple” and “dipole” for the allowed lines. Under the earlier definition a simple line has one hull intersection, whereas a dipole has two boundary intersections. If dipoles are intended to be included, the displayed example applies. If only single-root lines are intended, the connectedness argument of Lemma 5.2 extends to a general continuum, with no fullness conclusion. Those are distinct valid and invalid scopes.

A reliable way to formulate a peripheral rigidity theorem is to specify connectedness of the boundary pieces and of their open interiors, as in Definition 5.5. Calling the local corners perfect does not remove the need to prove the global connectedness of the intersection. Once those hypotheses are established, Proposition 7.2 applies to J itself, even when J has complementary holes.

9.2 A precise null-attachment model

We now give a complete metric formulation of the body-and-limbs geometry. It separates what follows from a null family from what requires additional access or puzzle construction.

Definition 9.2 (Attached null family)**.** Let B be a closed Jordan disk. An attached null family consists of nondegenerate continua L_j and distinct roots $a_j \in \partial B$ such that

$$L_j \cap B = \{a_j\}, \qquad L_i \cap L_j = \varnothing \ (i \neq j), \qquad \operatorname{diam} L_j \longrightarrow 0.$$

Set $K = B \cup \bigcup_{j\geq 1} L_j$. When exterior access statements are used, we additionally require K to be a hull. Finite families are included by omitting the limit condition.

Disjointness here is stronger than merely disjoint relative interiors. It reflects the distinct-root sector decomposition. More general trees of attachments can be treated recursively, but their intersections must be specified rather than inferred from a drawing.

Lemma 9.3 (Compactness and isolated off-body pieces)**.** *The set K in Definition 9.2 is compact and connected. Each $L_j \setminus \{a_j\}$ is relatively open in K. A sequence of points from pairwise different limbs can accumulate only on B.*

Proof. The family is bounded: all roots lie in the bounded set B, the diameters are eventually at most 1, and there are only finitely many exceptions. For a sequence $x_n \in L_{j_n}$ with distinct indices,

$$\operatorname{dist}(x_n, B) \leq |x_n - a_{j_n}| \leq \operatorname{diam} L_{j_n} \longrightarrow 0.$$

Every limit therefore lies in the closed set B. Sequences using only finitely many limbs have their limits in the finite union of those compact limbs and B. This proves closedness and compactness. Every limb is connected and meets the connected body, so their union is connected.

If $x \in L_j \setminus \{a_j\}$, its distance from B is positive. All but finitely many other limbs lie in a sufficiently small neighborhood of B and hence stay away from x. The remaining finitely many are compact and disjoint from L_j, so a smaller neighborhood of x misses them too. It intersects K only in $L_j \setminus \{a_j\}$, proving relative openness.

An explicit radius makes this isolation argument transparent. Put $d = \operatorname{dist}(x, B) > 0$, and choose N so that $\operatorname{diam} L_i < d/4$ for all $i > N$. For those indices and all $y \in L_i$,

$$|x - y| \geq \operatorname{dist}(x, B) - \operatorname{dist}(y, B) > 3d/4.$$

For each of the finitely many $i \leq N$ different from j, the compact set L_i misses x, so $d_i = \operatorname{dist}(x, L_i) > 0$. Choose $r < \min\{d/2,\, d_i/2 : i \leq N,\ i \neq j\}$, omitting an empty finite list. The ball $B(x, r)$ misses B and every other limb. This radius may depend strongly on x and tends to be uninformative as x approaches the root; openness at a root requires the different saturation argument. □

Proposition 9.4 (Local connectivity away from attachment roots)**.** *For every $a \in B \setminus \{a_j : j \geq 1\}$, the attached null continuum K is locally connected at a, whether or not its individual limbs are locally connected.*

Proof. Fix $\varepsilon > 0$. There are finitely many limbs with diameter at least $\varepsilon/3$. Since none has root a and all are disjoint from B except at their roots, choose a connected relatively open neighborhood W of a in B, contained in $B(a, \varepsilon/3)$, and avoiding the roots of those finitely many limbs. Such neighborhoods exist because a closed Jordan disk is locally connected. Form its saturation

$$O = W \cup \bigcup_{a_j \in W} L_j.$$

Every attached limb meets W at its root, so O is connected. Every selected limb has diameter less than $\varepsilon/3$, whence $O \subset B(a, 2\varepsilon/3)$.

The outer estimate applies to every point, not just roots. If $y \in L_j$ is selected, its root belongs to W, and hence

$$|y - a| \leq |y - a_j| + |a_j - a| \leq \operatorname{diam} L_j + |a_j - a| < 2\varepsilon/3.$$

For points of W the sharper bound $\varepsilon/3$ holds. Connectedness can likewise be checked for arbitrary many limbs: any separation of O would put the connected set W wholly on one side, and every connected L_j selected by $a_j \in W$ must lie on that same side because it meets W. The other side would be empty. Thus no locally connected structure inside an individual limb is needed when it is included whole.

It remains to prove that O is open in K. At a point off B this follows from Lemma 9.3: the entire corresponding limb is included. At a point $b \in W$, suppose points of $K \setminus O$ approach b. They cannot eventually belong to $B \setminus W$, since W is open in B. If their limb indices stay in a finite set, compactness puts b in an excluded limb, so b equals its root, contrary to $b \in W$. If the indices are distinct, the limb diameters tend to zero; their roots then approach b and eventually lie in W, again a contradiction.

For the last assertion, along a subsequence with distinct indices one has the explicit root estimate

$$|a_{j_n} - b| \leq |a_{j_n} - x_n| + |x_n - b| \leq \operatorname{diam} L_{j_n} + |x_n - b| \longrightarrow 0.$$

Because the roots lie in B and W is open in B, convergence to $b \in W$ forces eventual membership in W. Those limbs were included in O by its definition, contradicting $x_n \notin O$. Every sequence of indices has either a constant subsequence from a finite list or a subsequence of distinct indices, so these cases exhaust a possible failure of openness. In the metric space K, absence of a sequence from the complement converging to b is equivalent to the existence of a relative ball around b contained in O. Thus O is relatively open. The connected open sets so obtained form a neighborhood basis at a. □

This proof is genuinely local. A small non-locally-connected limb may remain non-locally-connected at its own off-body points; those defects are not repaired by attaching it to a Jordan body. At a non-root body point, however, each sufficiently small neighboring limb can be included whole in a connected neighborhood. No internal parametrization of that limb is required.

Proposition 9.5 (Locally connected gluing)**.** *If every limb in an attached null family is locally connected, then K is locally connected everywhere.*

Proof. Off B, local connectivity follows from relative openness of the corresponding limb. At non-root body points use the preceding proposition. At a root a_j, choose a small connected open neighborhood W_B in B and a small connected open neighborhood W_j in L_j, both containing a_j. Choose W_B to avoid all other roots of limbs whose diameters are larger than a prescribed small tolerance. The set

$$O = W_B \cup W_j \cup \bigcup_{i \neq j,\ a_i \in W_B} L_i$$

is connected and has small diameter. The same sequence argument as above proves openness at points of W_B other than possibly a_j. At a_j, points in the distinguished limb approaching it lie eventually in W_j, while points from other limbs are handled by their roots and null diameters. Openness elsewhere follows from relative openness of the limbs. This gives the required connected open neighborhoods at the remaining roots.

To keep the diameter tolerance explicit, first choose W_B and W_j inside $B(a_j, \varepsilon/4)$ and avoid the other finitely many roots whose limbs have diameter at least $\varepsilon/4$. Each remaining limb

included through W_B then lies in $B(a_j, \varepsilon/2)$, so

$$a_j \in O \subset B(a_j, \varepsilon/2), \qquad \operatorname{diam} O \le \varepsilon.$$

The distinguished limb is not included whole; its own connected open neighborhood W_j controls its size. All other selected limbs are included whole and controlled by the null-family bound. This division of roles explains why local connectivity of L_j is needed at its root, although the preceding proposition required no such hypothesis at a nonroot body point. The estimates and openness proof use the same fixed tolerance. □

9.3 Accessibility and the role of the boundary continuum

Assume now that the glued set is a locally connected hull. The exterior Caratheodory map lands at every boundary point. If $a \in \partial B \setminus \{a_j\}$, then $K \setminus \{a\}$ is connected: $B \setminus \{a\}$ is connected and every limb is still attached at its unchanged root. Two different accesses at a would give a ray-pair cut-line. Lemma 5.2 would then put nonempty parts of $K \setminus \{a\}$ in its two disjoint open sectors, a contradiction. Thus there is exactly one access and exactly one landing external angle.

This supplies a complete accessibility and uniqueness result in the locally connected gluing scope. Proposition 9.4 proves local connectivity at non-root body points even without that global assumption. Deriving every access and perfectly rigid boundary-piece conclusion of source Lemma 9.19 from its weaker hypotheses requires its additional geometric construction. Those conclusions are not being silently inferred from the local-connectivity proposition alone.

The distinction between K and $J = \partial K$ is particularly sharp here. A dipole cut for J may run through the interior of B; a finite cut-curve for K cannot. Proposition 7.3 therefore prevents hull-cut rigidity at points of ∂B, even in the simplest disk example. The boundary formulation is the appropriate place for such peripheral dipole puzzles.

9.4 Dense roots and a null-family criterion

Source correspondence. Lyubich, Lemma 9.20.

The following criterion makes the sector hypotheses explicit. It is useful when the small diameters are to be proved rather than assumed.

Proposition 9.6 (Dense sector attachments shrink)**.** *Let K be a hull containing a closed Jordan disk B. Suppose*

$$K = B \cup \bigcup_{j \ge 1} L_j, \qquad L_j \cap B = \{a_j\},$$

where the L_j are compact connected sets with distinct roots dense in ∂B. Assume each $L_j \setminus \{a_j\}$ is contained in an open sector S_j bounded by two external rays of K landing at a_j, that $S_j \cap B = \varnothing$, and that the sectors are pairwise disjoint. Then (L_j) is a null family.

Proof. Suppose infinitely many limbs have diameter at least $\varepsilon > 0$. Compactness of the hyperspace of nonempty compact subsets of K gives a Hausdorff convergent subsequence, with limit L_∞. It is a continuum of diameter at least ε. For each fixed j, all sufficiently late members of this subsequence miss the open sector S_j. Thus their limit misses every S_j. By the decomposition it is contained in B. The limbs miss $\operatorname{int} B$, so $L_\infty \subset \partial B$.

Two elementary Hausdorff estimates justify the retained size and location. For compact nonempty sets A, C,

$$|\operatorname{diam} A - \operatorname{diam} C| \le 2 d_H(A, C),$$

because approximating each endpoint of a pair changes its distance by at most twice the Hausdorff error. Thus the limit still has diameter at least ε. If a point $p \in L_\infty$ belonged to an open sector S_j, choose $r > 0$ with $B(p, r) \subset S_j$. Hausdorff convergence would put points of every sufficiently late limb in that ball, although those limbs miss S_j. The same open-ball argument excludes $p \in \operatorname{int} B$. Hence neither positive diameter nor the boundary location is merely inferred from the drawing of accumulating limbs; both follow from the definition of Hausdorff convergence.

A nondegenerate connected compact subset of a Jordan circle contains a nontrivial circle arc. Choose a root a_j in the interior of such an arc, and a different root a_k. Select one external ray at each root, and join the roots by a simple crosscut of B, whose interior lies in $\operatorname{int} B$. The two external rays and this crosscut form a proper embedded line M. It meets ∂B only at a_j, a_k and places the two open circle arcs between these roots on opposite sides. This last assertion is the Jordan separation theorem, after using the Schoenflies coordinate of B to straighten its boundary locally at the two endpoints.

Every L_n with $n \notin \{j, k\}$ misses M. Its only point in B is its own root, which is not an endpoint of the crosscut, and external rays miss K except at their landing points. Since L_n is connected, it lies in one open side of M. Pass to a further subsequence on which that side is fixed. Its Hausdorff limit lies in the corresponding closed side. But L_∞ contains a circle arc through a_j with a_j in its interior, and this arc has points on both open sides of M. This is the required contradiction. □

The proof uses two distinct roots to construct a genuine separating line through the body. Merely saying that a limit arc "crosses a limb" is not a substitute for that separation argument. Density supplies a root in the interior of the limit arc; compactness supplies the limit continuum; pairwise disjoint open sectors force that limit onto the body boundary. Each hypothesis has a separate role.

9.5 Solved investigations

Investigation 9.7. Compute the two components in Example 9.1 using angular intervals.

Solution. With angles modulo 2π, the components are $\{e^{it} : -\pi/6 \le t \le \pi/6\}$ and $\{e^{it} : 5\pi/6 \le t \le 7\pi/6\}$. They are nonempty disjoint compact arcs. Removing their endpoints gives the corresponding two components of $J \cap S$. Thus neither the closed nor the open middle intersection is connected.

Investigation 9.8. For an attached null family and $\eta > 0$, show that only finitely many limbs can contain a point at distance at least η from the body.

Solution. If $x \in L_j$ has $\operatorname{dist}(x, B) \ge \eta$, then $\operatorname{diam} L_j \ge |x - a_j| \ge \eta$. Nullness makes the set of such indices finite. This converts the global diameter condition into the local fact needed to isolate points away from the body.

Investigation 9.9. Can the conclusion of Proposition 9.4 imply that every limb is locally connected?

Solution. No. Attach a single non-locally-connected full continuum to a disk at one point, positioning the rest of it outside the disk. At a body point other than the attachment, small neighborhoods miss the limb entirely. The body-point conclusion holds there, while the limb retains its off-body defects. The proposition deliberately states a pointwise conclusion.

Investigation 9.10. Show why compactness and disjointness alone cannot replace the density and sector hypotheses in Proposition 9.6.

Solution. In a square, take disjoint vertical segments of length $1/2$ rooted at points of its bottom edge with abscissae tending to 0. Their Hausdorff limit is another segment of length $1/2$, not a point. The roots are not dense in a Jordan boundary, and the limiting segment is not forced onto that boundary by a decomposition into the required sectors. The example isolates the missing step in a bare compactness argument.

Investigation 9.11. Under locally connected gluing, why does a non-root point of the body boundary have exactly one access rather than merely at most one?

Solution. Connectedness of $K \setminus \{a\}$ excludes two accesses by the ray-pair cut argument, so gives only an upper bound. Existence is supplied separately by the continuous surjective exterior Caratheodory map for the locally connected hull. A preimage angle gives one landing ray. Combining the existence statement with the upper bound yields exactly one.

Chapter 10

Regulated legal arcs and finite trees

10.1 Pointing the interior components

Source correspondence. Lyubich, §9.3 and Exercises 9.23–9.26.

Let K be a locally connected hull. Choose a center b_U in each component U of $\operatorname{int} K$, and a Riemann map $\psi_U : \mathbb{D} \to U$ with $\psi_U(0) = b_U$. Its continuous extension is a homeomorphism by Theorem 8.3. An internal ray is

$$\rho_{U,\theta} = \{\psi_U(re^{i\theta}) : 0 \le r \le 1\}.$$

Changing the normalized Riemann map by a rotation only relabels the rays. Thus the set of internal rays depends on the chosen center, not on a remaining angular normalization. A legal point is a point of ∂K or one of the chosen centers.

For two boundary points of U, their radial connection is the union of their two internal rays through b_U. For a boundary point and the center, it is one internal ray. We call these the canonical radial arcs in $\overline{U}$. They are embedded arcs because ψ_U is a homeomorphism of the closed disk.

Definition 10.1 (Closed-component legal convention)**.** An arc α between legal points of K is legal if, for every interior component U, its intersection with $\overline{U}$ is empty, a singleton, or one canonical radial arc. The degenerate arc from a point to itself is also allowed. This is the regulated, closed-component interpretation of legal arcs used throughout this book.

The closure in this definition is necessary. If one constrains only $\alpha \cap U$, both boundary semicircles of a closed disk satisfy the condition vacuously, since they miss its interior. The diameter joining the same endpoints also consists of two internal rays. Uniqueness would therefore fail even for a single disk. The closed-component convention excludes the boundary semicircles and selects the radial connection. It states the normalization required for source Exercise 9.23 instead of using an ambiguous open-slice condition in a uniqueness proof.

Theorem 10.2 (Unique legal joining arc)**.** *Every pair of legal points $x, y \in K$ is joined by a unique legal arc, denoted $[x, y]_K$.*

Proof. Assume $x \neq y$, and choose an embedded arc $\gamma : [0, 1] \to K$ joining them. For any component U, the parameter set $\gamma^{-1}(\overline{U})$ is a closed interval, a singleton or empty. Indeed, if there were a complementary parameter interval with two endpoints mapping to $\overline{U}$, the intervening subarc would lie outside $\overline{U}$. Its gate under g_U would be constant, including both endpoint values by continuity. The endpoints would be the same point, contradicting injectivity of γ.

In parameter form, if $s < t$ are two contacts and a gap lies between them, take a component (c, d) of $[s, t] \setminus \gamma^{-1}(\overline{U})$. Closedness of the contact set gives $\gamma(c), \gamma(d) \in \overline{U}$. Gate constancy on the

open subarc and continuity at its endpoints give

$$\gamma(c) = g_U(\gamma(c)) = g_U(\gamma(d)) = \gamma(d), \qquad c < d,$$

contrary to the arc parametrization. Thus between any two contact parameters every intermediate parameter is also a contact. A closed subset of an interval with this interval-convexity property is a closed interval, a singleton or empty, exactly the alternatives required for unambiguous replacement.

Whenever this interval is nondegenerate, replace its subarc by the canonical radial connection between its endpoints. An endpoint in U can occur only when it is x or y, hence is the chosen center; every other endpoint is on ∂U. Thus the required canonical connection is defined. Intervals associated with different components have disjoint interiors: their images can overlap only at the single common boundary point allowed by Corollary 8.2. There are countably many such intervals, since each has a rational number in its interior.

For example, on an interval $[s,t]$ with two distinct boundary endpoints $p = \psi_U(e^{i\alpha})$, $q = \psi_U(e^{i\beta})$, a concrete continuous replacement is obtained by setting $v = (r-s)/(t-s)$ and using

$$R_U(r) = \begin{cases} \psi_U((1-2v)e^{i\alpha}), & 0 \le v \le 1/2, \\ \psi_U((2v-1)e^{i\beta}), & 1/2 \le v \le 1. \end{cases}$$

Both formulas give b_U at $v = 1/2$, and they give p, q at the two ends. Distinct boundary arguments make the two radial segments disjoint except at the center, so this parametrization is injective. For a center endpoint only the corresponding single radial segment is needed. The parameter speeds carry no conformal or length claim; they provide a definite continuous gluing on the original interval.

Make the replacements continuously on their original parameter intervals, fixing the endpoint values. A finite collection gives a continuous curve. The full sequence converges uniformly: on an as-yet-unreplaced interval both the original value and the replacement lie in $\overline{U}$, so

$$\sup_{t \text{ in that interval}} |\gamma(t) - \widetilde{\gamma}(t)| \le \operatorname{diam} U.$$

The interior components are a null family by Theorem 3.5. Consequently the supremum over all omitted intervals tends to zero, and the complete replacement is continuous.

To make this limit independent of an informal order of replacement, enumerate the nondegenerate intervals and let γ_N replace the first N. For a finite list keep the final curve unchanged thereafter, and define the supremum of an empty omitted list to be zero. The null-family property implies

$$d_N = \sup\{\operatorname{diam} U_j : j > N\} \longrightarrow 0, \qquad \|\gamma_M - \gamma_N\|_\infty \le d_N \quad (M > N).$$

Indeed for every $\eta > 0$ there are only finitely many involved components with diameter at least η, and every enumeration has included all of them by some finite index. At a parameter where the two curves differ, both values lie in the closure of the single component being replaced there. The bound is a supremum, not a sum of diameters; convergence does not require $\sum_j \operatorname{diam} U_j < \infty$. Completeness of the plane gives a uniform limit, and the usual three-term estimate with one continuous γ_N proves its continuity at every parameter, including an accumulation of replacement intervals. It is injective. New interior points lie inside their own component and cannot meet another replacement or the unchanged part of the original arc. Its boundary endpoints retain their original distinct parameter values. Thus the result is an embedded arc. Its intersection with each closed component is empty, the prescribed radial connection, or a retained singleton contact, so it is legal. An original singleton contact with a different component may disappear when a boundary subarc is replaced; this does not affect legality.

Uniform limits of injective curves need not be injective, so the separate geometric check above is necessary. If a final image value is an interior point of some U, it can arise only from that

component's replacement: the original contact interval contained every old visit to $\overline{U}$, and other components have disjoint interiors. Its parameter is therefore unique by injectivity of the radial parametrization. If the final value lies in ∂K, it is either an unchanged original value or an endpoint value retained at its original parameter. Thus

$$\widetilde{\gamma}(s) = \widetilde{\gamma}(t) \in \partial K \quad \Longrightarrow \quad \gamma(s) = \gamma(t) \quad \Longrightarrow \quad s = t.$$

The interior and boundary cases exhaust the image. This establishes an embedded arc after the infinite replacement, not merely a continuous path.

For uniqueness, suppose legal arcs α and β with the same endpoints differ. A component of $\alpha \setminus \beta$, together with the subarc of β between its two endpoint contacts, contains a Jordan curve Γ. By fullness its bounded complementary domain V lies in K, hence in one component U of $\operatorname{int} K$. Therefore $\Gamma \subset \overline{U}$. But each legal arc meets $\overline{U}$ in at most two internal rays, so Γ would be contained in at most four rays through b_U. Such a finite radial star contains no Jordan curve: removing its center leaves disjoint intervals, and two points in different intervals can be connected only through that center. This contradiction proves uniqueness.

The radial-star obstruction can also be checked in the disk coordinate. For a finite set of distinct angles Θ, put

$$T_\Theta = \bigcup_{\theta \in \Theta} [0, e^{i\theta}], \qquad T_\Theta \setminus \{0\} = \bigsqcup_{\theta \in \Theta} (0, e^{i\theta}].$$

Each component on the right is an interval. A Jordan curve avoiding 0 would be connected and contained in one such interval, which is impossible. A Jordan curve containing 0 would remain connected after that one point is removed, so its remainder would also have to lie in one radial interval; its closure still cannot form a circle. Pulling this conclusion forward by ψ_U excludes the alleged Γ. The argument works for any finite number of rays, explaining both the two-arc uniqueness proof and the later finite-hull no-cycle argument. □

The proof separates three inputs which serve different purposes: gates make each replacement interval unambiguous; the null-family theorem makes infinitely many replacements continuous; fullness excludes a second legal arc by filling a Jordan loop. A proof based only on the visual shape of a finite tree would not establish the existence statement for a hull with infinitely many interior components.

10.2 Boundary points in the interior of a legal arc

Proposition 10.3 (Legal interior points are cutpoints)**.** *Let $\alpha = [x, y]_K$ be a legal arc, and let $a \in \partial K \cap (\alpha \setminus \{x, y\})$. Then $K \setminus \{a\}$ is disconnected. Moreover there is a pair of external rays landing at a whose cut-line separates x and y. Thus a is accessible from the two sides of the arc.*

Proof. If x and y were in one component of $K \setminus \{a\}$, the local path connectivity of this open subspace would give an arc between them avoiding a. Apply the regulation construction from the preceding proof. It cannot introduce a: each inserted radial connection has all its new points in an interior component, while $a \in \partial K$, and its boundary endpoints were already on the original arc. We would obtain a legal arc from x to y avoiding a, contrary to uniqueness. Thus x and y lie in different components C_x, C_y of $K \setminus \{a\}$.

Each of these components contains a point of $\partial K \setminus \{a\}$. This is immediate if the chosen point already lies on the boundary. If it is in an interior component U, use an internal ray ending at a boundary point different from a; the Jordan boundary of U has more than one point, and that ray avoids a. Choose boundary points $x' \in C_x$, $y' \in C_y$ and angles θ_x, θ_y landing there under the continuous exterior map $\Phi : \mathbb{T} \to \partial K$.

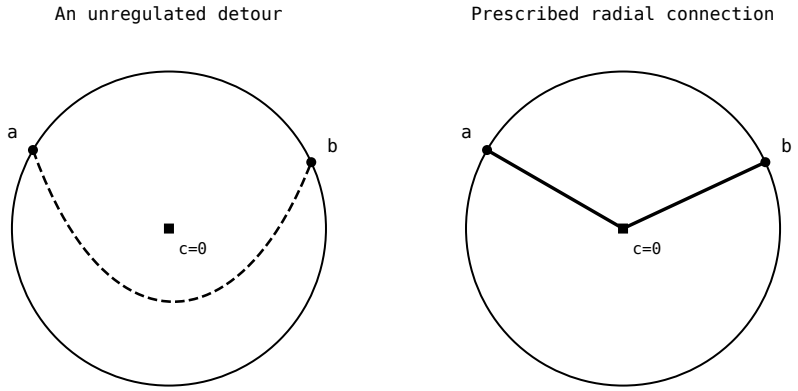


Figure 10.1: In one closed disk component, a curved boundary-to-boundary subarc is replaced by the two radial segments through the prescribed center. Endpoints are fixed. The theorem controls infinitely many such replacements by the null-family bound, not by the finite diagram or by a summable-length assumption.

Write $Q = \Phi^{-1}(a)$. Each of the two open circle arcs between θ_x and θ_y meets Q. Otherwise the image of the closed arc avoiding Q would be a connected subset of $K \setminus \{a\}$ joining x' and y'. Choose $\zeta_1, \zeta_2 \in Q$, one in each arc. Their external rays land at a and form a simple cut-line. The angles θ_x, θ_y lie in different shadows, so their landing points x', y' lie in opposite sectors. A connected component of $K \setminus \{a\}$ cannot cross that line; hence x and y are in opposite sectors too. Locally the two branches of the legal arc and the two external accesses alternate around a, which is the precise two-sided access statement.

The chosen angles have a definite alternating order. After choosing a starting point and lifting cyclic order to a real interval, reuse the same symbols for those real argument lifts and write

$$\theta_x < \zeta_1 < \theta_y < \zeta_2 < \theta_x + 2\pi.$$

The two components of the circle minus $\{\zeta_1, \zeta_2\}$ therefore contain θ_x and θ_y separately. Uniformization transfers these two angular regions to the two sides of the ray pair. Since the only point of that pair lying in K is a, each connected component of $K \setminus \{a\}$ lies entirely in one side. The separation of x', y' therefore also separates the original x, y. This order selection is what ensures that the ray pair detects the specified legal branches, even when the full landing fiber has more than two angles. □

The last argument uses a complete circle fiber, not merely two arbitrarily chosen landing angles. A cutpoint may have many accesses. Selecting one angle in each of the two separating circle intervals ensures that the resulting ray pair separates the particular branches of the legal arc under consideration.

Proposition 10.4 (Finite landing fibers and branches)**.** *Let $a \in \partial K$ have exactly $m \geq 2$ landing angles under the exterior Caratheodory map. In their cyclic order, each open sector between consecutive rays has connected intersection with K. Its closure in K is a branch rooted at a. Thus a is well branched with valence m.*

Proof. Let I be the open circle interval between two consecutive angles of $Q = \Phi^{-1}(a)$, and S the corresponding sector. The continuous image $\Phi(I)$ is connected and does not contain a. A ray with angle in I lies in S and can land only in $\overline{S}$; its landing point is not a, so lies in S.

Conversely every point of $\partial K \cap S$ has a preimage angle, and its landing ray cannot cross either boundary ray. Its angle therefore belongs to I. Hence

$$\partial K \cap S = \Phi(I)$$

is connected.

Every interior component meeting S lies wholly in S, since its interior misses the bounding cut-line. Its Jordan boundary has a point in S: it cannot be the singleton a. The union of that component with its boundary portion in S is connected and meets $\Phi(I)$. Adding all such components preserves connectedness and gives $K \cap S$. The one-cut closure assertion then gives $\overline{K \cap S} = (K \cap S) \cup \{a\}$. Doing this for all m consecutive angular intervals gives the branches. □

This proves a precise locally connected scope of source Corollary 9.13. The restriction $m \geq 2$ is part of the conclusion's meaning: an endpoint of a simple arc has one access and is not a cutpoint of valence at least two. Extending the argument to a merely rigid point in a globally non-locally-connected hull requires a substitute for the continuous circle boundary map; that stronger source scope is recorded separately.

10.3 Symmetry includes the centers

Proposition 10.5 (Equivariant pointing)**.** *Suppose $K = -K$ and the chosen centers satisfy $b_{-U} = -b_U$ for every interior component. Then for any legal point x, the legal arc $[-x, x]_K$ contains* 0.

Proof. The map $z \mapsto -z$ sends each pointed component (U, b_U) conformally to $(-U, b_{-U})$. Composing its Riemann maps gives a disk automorphism fixing 0, hence a rotation. It therefore sends internal rays to internal rays and legal arcs to legal arcs. Uniqueness yields

$$-[-x, x]_K = [-x, x]_K.$$

If $x = 0$ there is nothing to prove. Otherwise parametrize this arc by a homeomorphism $\alpha : [0, 1] \to [-x, x]_K$. Negation induces a continuous involution h of $[0, 1]$ which exchanges its endpoints, hence is decreasing. The function $h(t) - t$ is strictly decreasing, positive at 0 and negative at 1. It has one zero t_*. At that point $\alpha(t_*) = -\alpha(t_*)$, so $\alpha(t_*) = 0$. □

Example 10.6 (Geometric symmetry without pointed symmetry)**.** Let $K = \overline{\mathbb{D}}$ and choose $b = it$, with $0 < t < 1$, as its center. The hull and the legal endpoints $-1, 1$ are geometrically symmetric, but the pointing is not. Under

$$A_b(z) = \frac{z - it}{1 + itz}, \qquad A_b(0) = -it, \qquad A_b(\pm 1) = \frac{\pm 1 - it}{1 \pm it},$$

the two endpoint images have nonzero real parts, whereas $A_b(0)$ is strictly negative imaginary. Thus $A_b(0)$ lies on neither radial segment from 0 to $A_b(1)$ or $A_b(-1)$. The legal arc through b consequently misses the physical origin. Source Exercise 9.25 is therefore valid with symmetry of the pointed structure; geometric symmetry alone does not supply its conclusion.

The nonzero real parts in this counterexample can be checked without a geometric inference from the fractional expression. Multiplication by the conjugate denominator gives

$$A_b(1) = \frac{1 - t^2 - 2it}{1 + t^2}, \qquad A_b(-1) = \frac{-(1 - t^2) - 2it}{1 + t^2}.$$

Since $0 < t < 1$, the first real part is positive and the second negative. Every nonzero point on the corresponding radial segment has the same strict sign of real part as its endpoint. The point $A_b(0) = -it$ is nonzero and has real part zero, so it belongs to neither segment. Thus the example depends on the chosen center, not on a failure of geometric symmetry of the disk or on a numerical approximation of the rays.

10.4 The finite legal hull

For a finite set $X = \{x_1, \dots, x_n\}$ of legal points, define

$$H(X) = \bigcup_{i,j} [x_i, x_j]_K.$$

This is a finite union of compact arcs and is compact. Its shape may have curved edges and may pass through interior centers, but its graph structure is finite.

Theorem 10.7 (Finite legal hulls are trees)**.** *For $n \geq 2$, $H(X)$ is a finite topological tree. It equals $\bigcup_{j=2}^{n} [x_1, x_j]_K$, and is the smallest connected union of legal arcs containing X. Its leaves belong to X; after suppressing vertices of degree two, it has at most $n-2$ branch vertices.*

Proof. A finite union of legal arcs contains no Jordan curve. Indeed any such curve would bound a domain in $\operatorname{int} K$, hence lie in the closure of one interior component. Its intersection with the finitely many legal arcs would be contained in a finite internal radial star, which has no Jordan curve.

Build $T_m = \bigcup_{j=2}^{m} [x_1, x_j]_K$ inductively. The initial T_2 is an arc. Suppose T_m is a finite tree and add $\alpha = [x_1, x_{m+1}]_K$. The intersection $\alpha \cap T_m$ is connected. Otherwise a component of $\alpha \setminus T_m$ between two intersection points, together with the unique path in T_m joining those points, would contain a Jordan curve. Since the intersection contains x_1, it is an initial closed subarc of α. The remaining tail, if nonempty, is attached at exactly one point of T_m. The result is again a finite tree. This also proves that its leaves are among the listed endpoints.

The induction adds at most one new tail at a time. If the attachment point is in the interior of an existing edge, subdivide that edge there; otherwise it is already a vertex. Subdivision changes (V, E) by $(1, 1)$, and adjoining a nonempty tail with its new endpoint also changes them by $(1, 1)$. Thus

$$E - V = -1$$

is preserved from the initial single edge. If the new endpoint already lies in the old tree, the tail is empty and no edge is added. Old leaves can disappear by gaining another incident edge, but a new leaf can only be the newly listed endpoint. This explains finiteness and the leaf inclusion before the later degree-sum estimate; an arbitrary finite union of unrelated plane arcs would not have this controlled intersection property and could contain cycles or more complicated intersections.

The unique tree path between x_i and x_j is a legal arc. In every closed component its segments belong to the same radial star, and the path between its entry and exit is their unique radial connection; there can be no separate reentry without creating a cycle. By Theorem 10.2, this path is $[x_i, x_j]_K$. Thus all pairwise arcs already lie in T_n, proving the displayed equality and minimality.

For the branch count, let ℓ be the number of leaves and b the number of vertices of degree at least three after suppressing degree-two vertices. The degree sum and the tree identity $E = V - 1$ give

$$\sum_{v} (\deg v - 2) = -2, \qquad \sum_{\deg v \geq 3} (\deg v - 2) = \ell - 2.$$

Each summand on the left is at least 1, so $b \leq \ell - 2 \leq n - 2$. This is a combinatorial bound, independent of Euclidean edge lengths or the complexity of the surrounding hull. □

The smallest legally convex set in the source terminology is obtained by including all these canonical connections. The closed-component convention is again essential: if boundary detours were legal, the union of all joining arcs in a disk could contain a circle instead of a tree. The quantitative chapters study finite trees after this topological construction, with a separately specified length or conductance model.

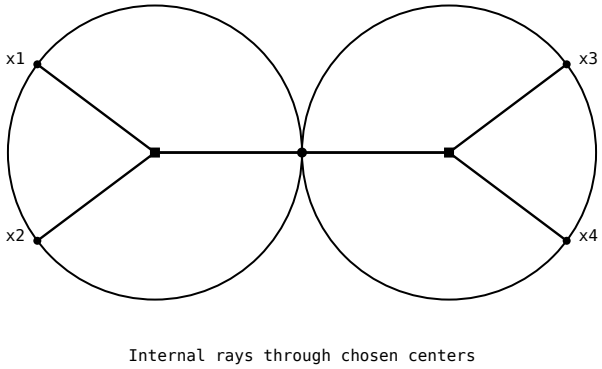


Figure 10.2: A finite legal hull uses internal radial connections in closed interior components. The resulting graph is a tree; the drawing's edge lengths carry no conformal or electrical normalization.

10.5 Solved investigations

Investigation 10.8. For the closed unit disk pointed at 0, list three distinct arcs from -1 to 1 allowed by an open-interior-only ray condition, and identify which survives Definition 10.1.

Solution. The upper semicircle, the lower semicircle and the diameter are distinct embedded arcs with the same endpoints. The semicircles have empty intersection with $\mathbb{D}$, so the open-interior condition imposes no restriction on them. The diameter is two internal rays. Under the closed-component convention their intersection with the one closed component is the entire arc; only the diameter is a canonical radial connection. Hence uniqueness holds exactly after the stated normalization.

Investigation 10.9. Let three distinct boundary points lie on one Jordan interior component. What is their legal hull, and what branch count does the theorem give?

Solution. It is the union of the three internal rays from the chosen center to those points. The center is the unique branch vertex and has degree three; there are three leaves. The bound gives $b \le 3 - 2 = 1$, attained here. Changing the center changes the geometric arcs but preserves this abstract three-leaf tree structure.

Investigation 10.10. If a component is fixed by $z \mapsto -z$, what does equivariance require of its chosen center?

Solution. It requires $b_U = -b_U$, hence $b_U = 0$. For components paired with a distinct reflected component, one may choose the center in one component freely and must choose its negative in the other. Thus symmetry of the pointing is concrete data, not merely a statement that the underlying hull is symmetric.

Investigation 10.11. A finite legal hull has five leaves and every branch vertex has degree three. How many branch vertices and edges remain after suppressing degree-two vertices?

Solution. The identity $\sum_{\deg v \ge 3}(\deg v - 2) = \ell - 2$ gives $b = 3$. There are $V = 5 + 3 = 8$ vertices and $E = V - 1 = 7$ edges. If a branch vertex had larger degree, fewer than three branch vertices could suffice; the number of leaves alone gives an upper bound until the degrees are specified.

Investigation 10.12. Why does the regulation construction require diameters tending to zero, rather than only the fact that the replacement intervals have lengths tending to zero in the parameter interval?

Solution. A very short parameter interval may be mapped across a large component. Replacing it by a curve with a distant excursion could destroy continuity at an accumulation of such intervals. The null-family estimate bounds the *image* displacement uniformly by $\operatorname{diam} U$, regardless of parameter speed. It yields uniform convergence of finite replacements and therefore continuity of the limit. Parameter lengths alone control no such image excursion.

Chapter 11

Closed identifications and the topology of a quotient

Source correspondence. Lyubich, §9.4, with the quotient-space prerequisites made explicit.

Write $T = \partial\mathbb{D}$. A pinched disk is obtained by identifying specified subsets of $\overline{\mathbb{D}}$, but the instruction to identify sets does not by itself define a usable topological model. Three questions have to be settled: do the proposed sets form a partition; is the relation closed; and what topology does the quotient inherit? These questions precede any assertion that the resulting space can be embedded in a plane.

Throughout this chapter X is a compact metric space. An equivalence relation R on X is a subset of $X \times X$ satisfying reflexivity, symmetry and transitivity. Its class through x is

$$[x]_R = \{y \in X : (x, y) \in R\}.$$

The quotient $Y = X/R$ is the set of classes, equipped with the topology in which $V \subset Y$ is open exactly when $q^{-1}(V)$ is open in X. Here $q(x) = [x]_R$. A set is *saturated* when it is a union of whole classes. In particular, $q^{-1}(q(A))$ is the saturation of A. We use the notation

$$R[A] = \{x \in X : \text{there is } a \in A \text{ with } xRa\}.$$

Even when all classes are closed, the relation itself need not be closed in $X \times X$. The distinction will matter in limiting constructions.

The algebra behind saturation will be used repeatedly. Symmetry and transitivity give, for arbitrary subsets $A, B \subset X$,

$$A \subset R[A], \qquad R[R[A]] = R[A], \qquad A \subset B \implies R[A] \subset R[B].$$

For example, if $x \in R[R[A]]$, choose $y \in R[A]$ with xRy, and then $a \in A$ with yRa. Transitivity gives xRa. The reverse inclusion follows from reflexivity. Consequently the identities $R[A] = A$ and $q^{-1}(q(A)) = A$ are equivalent ways of saying that A is saturated. A complement of a saturated set is saturated as well: a class cannot partly lie in that set and partly in its complement.

11.1 Closedness and upper semicontinuity

Theorem 11.1 (Closed-relation criterion)**.** *If R is closed in $X \times X$, then its classes are compact, the quotient Y is compact and Hausdorff, and q is a closed map. Moreover, if a class A is contained in an open set $U \subset X$, there is a saturated open set W such that*

$$A \subset W \subset U.$$

Conversely, if the quotient topology on Y is Hausdorff, then R is closed.

Proof. A class is the inverse image of R under the continuous map $y \mapsto (x, y)$, so it is closed and compact. If $F \subset X$ is closed, the set

$$R \cap (X \times F)$$

is compact. Its projection onto the first coordinate is $R[F]$, which is compact and therefore closed in the metric space X. Now $q^{-1}(q(F)) = R[F]$. The definition of the quotient topology says that $q(F)$ is closed precisely when its inverse image is closed. Thus q is closed.

Here the compactness step can be expressed as an exact projection:

$$R[F] = \pi_1\big(R \cap (X \times F)\big), \qquad \pi_1(x, y) = x.$$

Both factors in the intersection are closed in the compact space X^2. Continuous images of compact sets are compact, and compact subsets of a metric space are closed. Thus no assumption that an arbitrary projection is a closed map is being used. The compact domain is the reason this particular projection preserves closedness. This observation is also valid when F is empty, in which case both displayed sets are empty.

For a class $A \subset U$, put

$$W = X \setminus R[X \setminus U].$$

This set is open. It is saturated because transitivity makes $R[B]$ saturated for every B. A class meeting W cannot meet $X \setminus U$, so $W \subset U$. The entire class A lies in W, since $A \subset U$.

More explicitly, membership has the useful universal form

$$x \in W \quad \Longleftrightarrow \quad [x]_R \cap (X \setminus U) = \varnothing \quad \Longleftrightarrow \quad [x]_R \subset U.$$

It follows that W is the largest saturated subset of U, not just one possible choice. If $V \subset U$ is saturated and $x \in V$, then $[x]_R \subset V \subset U$, so $x \in W$. Once openness of W has been proved, compactness of A also yields a metric neighborhood of the entire class: unless $W = X$, the number $d(A, X \setminus W)$ is positive. Every sufficiently nearby representative therefore belongs to a class contained in U.

This proves the neighborhood assertion, including its quantifiers:

$$\forall A \in X/R \ \forall U \text{ open with } A \subset U \ \exists W \text{ open and saturated}: \quad A \subset W \subset U.$$

For distinct classes A, B, compactness and the metric give disjoint open neighborhoods U_A, U_B. Shrink them by the preceding construction to saturated open neighborhoods W_A, W_B. Their images under q are open by the quotient definition, and are disjoint because saturated disjoint sets cannot contain representatives of the same class. They separate $q(A)$ and $q(B)$. Compactness of Y follows from continuity and surjectivity of q.

To make the metric separation explicit, disjoint nonempty compact classes have

$$\rho = \min\{d(a, b) : a \in A,\ b \in B\} > 0.$$

The sets $\{x : d(x, A) < \rho/3\}$ and $\{x : d(x, B) < \rho/3\}$ are disjoint: a point in both would give $d(A, B) < 2\rho/3$ by the triangle inequality. After saturation shrinking, their quotient images are still disjoint because

$$q^{-1}(q(W_A)) = W_A, \qquad q^{-1}(q(W_B)) = W_B.$$

This is why merely projecting two ordinary disjoint open sets would not have been enough; they might meet different representatives of the same class.

Conversely, the diagonal Δ_Y of a Hausdorff space is closed. Since

$$R = (q \times q)^{-1}(\Delta_Y),$$

continuity of $q \times q$ proves that R is closed.

For the diagonal assertion, if $u \neq v$ in a Hausdorff space, choose disjoint open neighborhoods O_u, O_v. The product $O_u \times O_v$ misses Δ_Y. Such products cover the complement of the diagonal, so that complement is open. Pullback now gives the explicit open set

$$(X \times X) \setminus R = (q \times q)^{-1}((Y \times Y) \setminus \Delta_Y).$$

This direction does not depend on a metric on Y. The topology used on the quotient is essential: it makes q continuous and hence makes this pullback argument valid.

□

The neighborhood assertion is often called upper semicontinuity of the decomposition. It controls a whole nearby class, not just a chosen representative. A useful sequential form is

$$x_n R y_n, \qquad x_n \longrightarrow x, \qquad y_n \longrightarrow y \quad \Longrightarrow \quad xRy.$$

On a metric product this is exactly closedness of R. It allows classes to merge in a limit, but it does not allow a pair of related points to converge to unrelated points.

Proposition 11.2 (Fiber limits)**.** *Let $f : X \to Z$ be continuous, with Z Hausdorff, and let $A_n = f^{-1}(z_n)$ be nonempty fibers. If $z_n \to z$ and $a_{n_j} \in A_{n_j}$ converges to a, then $a \in f^{-1}(z)$. In particular, every Hausdorff limit of the sets A_n is contained in $f^{-1}(z)$.*

Proof. Continuity gives $f(a_{n_j}) \to f(a)$, while the assumed convergence gives $f(a_{n_j}) = z_{n_j} \to z$. Uniqueness of limits in a Hausdorff space gives $f(a) = z$. If A_{n_j} converges in Hausdorff distance to A, every $a \in A$ is the limit of points $a_{n_j} \in A_{n_j}$. Applying the first assertion pointwise gives $A \subset f^{-1}(z)$. Equality need not hold: new points may enter the limiting fiber without being approached by the selected earlier fibers.

There is also an eventual-containment consequence useful in practice. If $U \subset X$ is open and contains $f^{-1}(z)$, then

$$A_n \subset U \qquad \text{for all sufficiently large } n.$$

Otherwise choose infinitely many $a_n \in A_n \setminus U$. Compactness gives a convergent subsequence; its limit lies in the closed set $X \setminus U$ and, by the proved assertion, in $f^{-1}(z) \subset U$. The contradiction proves containment. This is an upper estimate on all points of nearby fibers. It supplies no reverse estimate forcing every point of the limiting fiber to be approximated.

□

This proposition is the precise content needed from the limiting-fiber step in the source proof of Theorem 9.28. It requires neither a differentiable boundary nor finite fibers.

11.2 Maps defined by classes

Proposition 11.3 (The factorization property)**.** *Let $q : X \to X/R$ be the quotient map and $g : X \to Z$ a continuous map. There is a continuous map $\bar{g} : X/R \to Z$ with $g = \bar{g} \circ q$ if and only if g is constant on every R-class. The factor is unique. If Z is Hausdorff, g is onto, and its fibers are exactly the R-classes, then $\bar{g}$ is a homeomorphism.*

Proof. Constancy on classes is necessary from the displayed factorization. If it holds, define $\bar{g}([x]) = g(x)$. This is well defined and unique because q is onto. For every open set $V \subset Z$,

$$q^{-1}(\bar{g}^{-1}(V)) = g^{-1}(V)$$

is open, so the quotient definition makes $\bar{g}^{-1}(V)$ open. Thus $\bar{g}$ is continuous. Under the last hypotheses it is a continuous bijection from a compact space to a Hausdorff space. The image of every closed subset is compact and hence closed, proving continuity of its inverse.

The two bijectivity tests are separate. Surjectivity follows from $\bar{g}(q(X)) = g(X) = Z$. For injectivity,

$$\bar{g}([x]) = \bar{g}([x']) \implies g(x) = g(x') \implies [x] = [x'].$$

The last implication is exactly the hypothesis that the fibers are the prescribed classes, rather than merely unions of such classes. In the homeomorphism case, a closed set C in the compact quotient has compact image $\bar{g}(C)$, which is closed in Z. Thus the inverse preimage of every closed set is closed. This verifies inverse continuity without choosing representatives continuously.

□

One immediate application is a boundary map $b : T \to \partial K$ that is continuous and onto. The relation

$$s \sim_b t \quad \iff \quad b(s) = b(t)$$

is closed, and the induced map $T/\!\sim_b \to \partial K$ is a homeomorphism. This statement concerns the boundary alone. It says nothing yet about filling the disk or extending the exterior map.

Proposition 11.4 (A countable quotient base)**.** *For a closed relation on a compact metric space, the quotient has a countable base.*

Proof. Fix a countable base $\mathcal{B}$ of X and let $\mathcal{F}$ be its countable family of finite unions. For $B \in \mathcal{F}$, set

$$W_B = X \setminus R[X \setminus B].$$

These sets are open and saturated. Let V be open in Y and let $y \in V$. The compact fiber $q^{-1}(y)$ lies in the open set $q^{-1}(V)$. Finitely many basic open sets cover the fiber and are contained in $q^{-1}(V)$; their union is a suitable B. Then

$$y \in q(W_B) \subset q(B) \subset V.$$

The sets $q(W_B)$ therefore form a countable base. The construction is an existence proof using all finite unions; it is not an algorithm for deciding membership in an arbitrary equivalence class.

One can see the countability without making a choice for every quotient point. Enumerate the original base as $(B_k)_{k\geq 1}$ and index finite unions by finite subsets of $\mathbb{N}$. The family of such subsets is countable because

$$\{I \subset \mathbb{N} : I \text{ finite}\} = \bigcup_{m\geq 0} \{I \subset \mathbb{N} : |I| = m\},$$

and each fixed-cardinality family is encoded by an increasing integer tuple. The same fixed countable family of saturated open sets works for all y, V. The finite covering argument establishes the base property after this family has been specified.

□

11.3 Connected fibers and connected images

Definition 11.5 (Monotone map)**.** A continuous surjection between compact Hausdorff spaces is *monotone* if every point inverse is connected. This usage has no connection with an order-preserving real function.

Lemma 11.6 (Connected inverse images)**.** *Let $q : X \to Y$ be a continuous closed surjection with connected fibers. For every connected subset $A \subset Y$, the set $q^{-1}(A)$ is connected. No closedness assumption on A is needed.*

Proof. The restricted map $q_A : q^{-1}(A) \to A$ is closed. Indeed, a relatively closed subset of $q^{-1}(A)$ has the form $F \cap q^{-1}(A)$ with F closed in X, and

$$q\big(F \cap q^{-1}(A)\big) = q(F) \cap A.$$

Suppose $q^{-1}(A) = U \cup V$ were a separation into disjoint nonempty relatively open and closed sets. A connected fiber lies wholly in one of U, V, so $q(U)$ and $q(V)$ are disjoint. They are nonempty, cover A, and are relatively closed because q_A is closed. They would separate A, a contradiction.

Indeed, if both U and V met the fiber over some a, their intersections with that fiber would be disjoint nonempty relatively open sets covering a connected space. Therefore each fiber lies on one side, and

$$q_A^{-1}(q(U)) = U, \qquad q_A^{-1}(q(V)) = V, \qquad A = q(U) \sqcup q(V).$$

Since the two images are relatively closed and complementary, they are relatively open as well. All of these statements are in the subspaces $q^{-1}(A)$ and A. In particular, there is no step requiring that A itself be closed in Y.

□

Theorem 11.7 (Local connectivity passes to compact images)**.** *Let X be compact and locally connected and let $f : X \to Y$ be a continuous surjection onto a Hausdorff space. Then Y is locally connected. The fibers need not be connected.*

Proof. Fix $y \in Y$ and an open neighborhood U of y. For every $x \in f^{-1}(y)$ choose a connected open neighborhood V_x of x contained in $f^{-1}(U)$. Compactness of the fiber supplies finitely many such neighborhoods covering it. Let their union be V. Every connected set $f(V_x)$ contains y, so

$$C = f(V) = \bigcup f(V_x)$$

is connected and lies in U. Since f is a closed map,

$$W = Y \setminus f(X \setminus V)$$

is an open neighborhood of y contained in C. Hence every open neighborhood of y contains a connected neighborhood of y.

The neighborhood inclusion can be checked one point at a time. If $y' \in W$, then no representative of the fiber $f^{-1}(y')$ belongs to $X \setminus V$. Hence

$$f^{-1}(y') \subset V, \qquad y' \in f(V) = C \subset U.$$

Also $y \notin f(X \setminus V)$ because the chosen finite cover contains its entire fiber. Thus $y \in W \subset C$. The connectedness of C uses the common image point y; it does not require that the neighborhoods V_x meet one another in X, and does not assume connected fibers.

For completeness, this yields connected *open* neighborhoods as well. Let G be open and let C_y be the component of G containing y. Every $z \in C_y$ has a connected neighborhood $N_z \subset G$. Its union with C_y is connected because they meet at z. Maximality of the component gives $N_z \subset C_y$. Each N_z contains an open neighborhood of z, so C_y is open. Components of open sets therefore give the required base. □

The compactness in this argument is structural. It makes the map closed and reduces an entire fiber to finitely many neighborhoods. In the disk model it will apply to the quotient of the closed disk after the quotient has been shown Hausdorff.

11.4 Worked investigations on quotient topology

Investigation 11.8 (A single collapsed compact set)**.** Let A be a nonempty compact subset of a compact metric space X. Identify all points of A and leave the other points separate. Prove that the quotient is Hausdorff. Decide when the quotient map is monotone.

Solution. The relation is

$$R = \Delta_X \cup (A \times A),$$

a union of two closed sets. Theorem 11.1 applies. Its only possibly non-singleton fiber is A, so the map is monotone exactly when A is connected. Hausdorffness alone imposes no connectedness on A. For example, identifying the two endpoints of $[0, 1]$ gives a circle, but the quotient map has a disconnected endpoint fiber.

Investigation 11.9 (Closed classes do not suffice)**.** On $[0, 1]$, declare $\{1/n, 1 - 1/n\}$ a class for each integer $n \geq 3$, and leave all other points separate. Is the relation closed?

Solution. For $n \geq 3$ the pairs are mutually disjoint: their first coordinates lie in $(0, 1/3]$, their second coordinates lie in $[2/3, 1)$, and each coordinate sequence is injective. Thus this is a partition into finite closed sets. But

$$(1/n, 1 - 1/n) \longrightarrow (0, 1),$$

and 0 and 1 belong to different singleton classes. The relation is not closed, so the quotient is not Hausdorff. Adding the class $\{0, 1\}$ repairs this particular defect; all possible off-diagonal limits then lie in that added class.

Investigation 11.10 (Two quotients of a circle)**.** Describe the quotients of T induced by $z \mapsto \operatorname{Re} z$ and by $z \mapsto z^2$. Which of the two quotient maps is monotone?

Solution. The first image is $[-1, 1]$, and its fibers are the conjugate pairs $\{e^{it}, e^{-it}\}$, with singleton fibers at ± 1. The second image is T, with antipodal fibers $\{z, -z\}$. Proposition 11.3 identifies the two quotient spaces with their stated images. Neither map is monotone: a two-point fiber is disconnected. Their later disk extensions behave very differently, because conjugate pairs are unlinked whereas distinct antipodal pairs have crossing convex hulls.

Investigation 11.11 (Disconnected fibers and local connectivity)**.** Explain why Theorem 11.7 does not imply that every continuous map from a continuum is monotone.

Solution. The conclusion of that theorem concerns neighborhoods in the image, not connectedness of point inverses. The map $t \mapsto e^{2\pi i t}$ from $[0, 1]$ onto T is a continuous surjection between compact locally connected continua, but the inverse of 1 is $\{0, 1\}$. The proof of the theorem deliberately takes images of several connected neighborhoods whose common intersection is only the image point. It never asserts that their union in the domain is connected.

Investigation 11.12 (A factorization test)**.** Let $A \subset X$ be compact and let q collapse A alone. For a continuous real function g on X, give a necessary and sufficient condition that g descend to the quotient. Apply this to $X = \overline{\mathbb{D}}$, $A = [-1, 1]$, and $g(z) = \operatorname{Re} z$ or $g(z) = \operatorname{Im} z$.

Solution. The condition is constancy on A, by Proposition 11.3. The real-part function is not constant on the diameter and cannot descend. The imaginary-part function is zero there and does descend. Its quotient function remains continuous even though the image of the diameter is a pinched point. This is a direct test of a proposed coordinate on a quotient; it cannot be replaced by drawing the coordinate contours after identification.

Chapter 12

Unlinked circle fibers and hyperbolic convex hulls

Source correspondence. Lyubich, Lemma 9.27 and the convex-hull construction preceding Theorem 9.28.

There are two geometries in the disk construction. Circular order detects whether two classes interlace. Hyperbolic geodesics then fill each class by a connected set. A Euclidean coordinate change separates these tasks: the circle order stays fixed while hyperbolic geodesics become straight chords. This makes both the topology and the limiting arguments explicit.

12.1 The Klein coordinate and finite convex combinations

Define

$$\kappa(z) = \frac{2z}{1+|z|^2}, \qquad z \in \overline{\mathbb{D}}.$$

The radial function $k(r) = 2r/(1+r^2)$ satisfies

$$k(0) = 0, \quad k(1) = 1, \quad k'(r) = \frac{2(1-r^2)}{(1+r^2)^2} > 0 \quad (0 \le r < 1).$$

Consequently κ is a homeomorphism of $\overline{\mathbb{D}}$ fixing T pointwise. Its inverse has the continuous expression

$$\kappa^{-1}(w) = \frac{w}{1+\sqrt{1-|w|^2}}, \qquad |w| \le 1.$$

The expression is valid at $w = 0$ as well as on T.

To derive the inverse, let $s = |w|$ and $r = |z|$. The equation $s = 2r/(1+r^2)$ gives

$$sr^2 - 2r + s = 0, \qquad r = \frac{1-\sqrt{1-s^2}}{s} = \frac{s}{1+\sqrt{1-s^2}} \quad (0 < s \le 1).$$

The other quadratic root is greater than one when $0 < s < 1$, so it is outside the required disk. The final expression tends to zero as $s \to 0$, and the angular coordinate is unchanged. This proves the inverse formula including its endpoint limits, rather than using only strict radial monotonicity to assert an unspecified inverse.

A circular geodesic orthogonal to T has center c and radius $\sqrt{|c|^2-1}$. Its equation is

$$|z-c|^2 = |c|^2 - 1 \quad \Longleftrightarrow \quad 2\,\mathrm{Re}(\bar{c}z) = 1 + |z|^2 \quad \Longleftrightarrow \quad \mathrm{Re}(\bar{c}\,\kappa(z)) = 1.$$

Thus κ sends it to the straight chord with the same endpoints. A diameter is already a geodesic and remains on the same line. This proves, without a metric approximation, that hyperbolic segments are carried to Euclidean segments.

Definition 12.1 (The closed ideal hull)**.** For a nonempty compact set $A \subset T$, define its closed hyperbolic convex hull by

$$H(A) = \kappa^{-1}\big(\operatorname{conv} A\big).$$

Here conv means the set of finite convex combinations in $\mathbb{R}^2 = \mathbb{C}$. A singleton has its singleton hull; two points have a closed geodesic as hull. An ideal polygon with interior requires at least three distinct boundary points.

The last distinction corrects a possible degeneracy in the source's earlier convex-hull terminology. A two-point hull is not a Jordan disk. It nevertheless is a legitimate non-separating continuum for a decomposition.

Lemma 12.2 (Three-point representation)**.** *Every point of* $\operatorname{conv} A$ *is a convex combination of at most three points of* A*. If* A *is compact, then* $\operatorname{conv} A$ *is compact and*

$$\operatorname{conv} A \cap T = A.$$

Proof. Start with a finite representation $z = \sum_{j=1}^{m} t_j a_j$, $t_j \geq 0$, $\sum t_j = 1$, and omit zero coefficients. If $m > 3$, the vectors $(a_j, 1) \in \mathbb{R}^3$ are linearly dependent. There are real s_j, not all zero, satisfying

$$\sum_j s_j a_j = 0, \qquad \sum_j s_j = 0.$$

Some s_j is positive and some is negative. Set $\tau = \min_{s_j>0} t_j/s_j$. Then $t'_j = t_j - \tau s_j$ are nonnegative, still sum to one and represent the same point; at least one is zero. Repeating reduces to at most three positive coefficients.

Every sign and normalization in this reduction can be checked directly:

$$\begin{aligned} s_j > 0 &\implies t_j - \tau s_j \geq 0, \\ s_j \leq 0 &\implies t_j - \tau s_j \geq t_j > 0, \\ \sum_j t'_j &= 1 - \tau \sum_j s_j = 1, \\ \sum_j t'_j a_j &= z - \tau \sum_j s_j a_j = z. \end{aligned}$$

The minimum defining τ is positive and is attained because the index set is finite and nonempty. At an attaining index the new coefficient is zero. Thus the number of positive coefficients strictly decreases, while the point and the total mass remain unchanged. After at most $m - 3$ repetitions the reduction stops.

It follows that $\operatorname{conv} A$ is the image of the compact set $A^3 \times \Delta_2$ under $(a_1, a_2, a_3; t_1, t_2, t_3) \mapsto \sum t_j a_j$, where $\Delta_2 = \{t_j \geq 0 : \sum t_j = 1\}$. Repeated points accommodate representations with fewer terms.

The parameter space is a finite product of compact metric spaces, and the barycenter map is continuous. A useful estimate making that continuity explicit is

$$\left| \sum_{j=1}^{3} t_j a_j - \sum_{j=1}^{3} s_j b_j \right| \leq \max_j |a_j - b_j| + \sum_{j=1}^{3} |t_j - s_j|.$$

To obtain it, add and subtract $\sum_j t_j b_j$, use $\sum_j t_j = 1$ for the first difference and $|b_j| = 1$ for the second. The map therefore sends convergent vertices and coefficients to convergent barycenters. Compactness of its image also proves that the set of finite convex combinations is already closed; no extra closure operation is missing from the definition.

Finally, let $\zeta \in T$ be such a convex combination. Taking the real part after multiplication by $\bar{\zeta}$ gives

$$1 = \sum_j t_j \operatorname{Re}(\bar{\zeta} a_j) \le \sum_j t_j = 1.$$

For every positive coefficient equality forces $\operatorname{Re}(\bar{\zeta} a_j) = 1$, hence $a_j = \zeta$. Thus $\zeta \in A$. The opposite inclusion is immediate.

Equivalently, the supporting-line calculation has the nonnegative defect identity

$$\sum_j t_j |a_j - \zeta|^2 = 2 \sum_j t_j (1 - \operatorname{Re}(\bar{\zeta} a_j)) = 0.$$

Each term is nonnegative, so every vertex of positive weight equals ζ. This deals with arbitrary zero weights and repeated vertices without any strict-convexity terminology. It also explains geometrically why adding convex combinations inside the disk cannot introduce a new ideal endpoint on its circle.

□

Proposition 12.3 (Continuity of the hull operation)**.** *For nonempty compact $A, B \subset T$,*

$$d_H(\operatorname{conv} A, \operatorname{conv} B) \le d_H(A, B).$$

If ω is a modulus of uniform continuity for κ^{-1} on $\overline{\mathbb{D}}$, then

$$d_H(H(A), H(B)) \le \omega\big(d_H(A, B)\big).$$

In particular, Hausdorff convergence of boundary sets implies Hausdorff convergence of their closed hyperbolic hulls.

Proof. Put $\delta = d_H(A, B)$. If $z = \sum_{j=1}^{3} t_j a_j$ belongs to $\operatorname{conv} A$, choose $b_j \in B$ with $|a_j - b_j| \le \delta$. The point $w = \sum_j t_j b_j$ belongs to $\operatorname{conv} B$ and

$$|z - w| \le \sum_j t_j |a_j - b_j| \le \delta.$$

Interchanging A and B gives the Hausdorff inequality. Applying uniform continuity of κ^{-1} to the selected pairs gives the second inequality. No claim that κ^{-1} is globally Lipschitz near T is required.

A concrete modulus is available from the inverse formula. For $u, v \in \overline{\mathbb{D}}$, write $s_u = \sqrt{1 - |u|^2}$ and $s_v = \sqrt{1 - |v|^2}$. Using $|\sqrt{a} - \sqrt{b}| \le \sqrt{|a - b|}$ for $a, b \ge 0$, one obtains

$$\begin{aligned} |\kappa^{-1}(u) - \kappa^{-1}(v)| &\le |u - v| + |s_u - s_v| \\ &\le |u - v| + \sqrt{\big||u|^2 - |v|^2\big|} \\ &\le |u - v| + \sqrt{2|u - v|}. \end{aligned}$$

The first inequality follows by subtracting the fractions and bounding both denominators below by one. The last uses $|u| + |v| \le 2$. Hence the theorem permits the explicit choice $\omega(t) = t + \sqrt{2t}$ on $0 \le t \le 2$. This is a uniform continuity bound, not a claim of Lipschitz regularity at the ideal boundary.

□

12.2 Circular order is the separation criterion

Definition 12.4 (Unlinked compact sets)**.** Two disjoint nonempty compact subsets $A, B \subset T$ are unlinked if there do not exist distinct $a_1, a_2 \in A$ and $b_1, b_2 \in B$ occurring in the alternating circular order

$$a_1,\ b_1,\ a_2,\ b_2.$$

A partition of T is unlinked if every two different classes are unlinked.

Lemma 12.5 (The chord criterion)**.** *For disjoint nonempty compact $A, B \subset T$, the following are equivalent:*

1. *A and B are unlinked;*
2. *B is contained in one component of $T \setminus A$;*
3. *$\operatorname{conv} A \cap \operatorname{conv} B = \varnothing$;*
4. *$H(A) \cap H(B) = \varnothing$.*

When A is a singleton, $T \setminus A$ is its single open-arc component and the same assertions hold.

Proof. For a singleton A, its convex hull is a point of T outside B, and Lemma 12.2 proves disjointness from $\operatorname{conv} B$. Assume henceforth that A has at least two points. If B is not contained in one component of $T \setminus A$, choose b_1, b_2 in two different components. Each of the two circle arcs joining b_1 to b_2 contains a point of A; otherwise that arc would join them inside $T \setminus A$. Choosing one such point on each arc gives an alternating quadruple. Conversely, such a quadruple prevents B from lying in a single complementary arc. This proves equivalence of the first two assertions.

Suppose B lies in a component I of $T \setminus A$. Its two distinct endpoints u, v lie in A. The line through u, v places I in one open half-plane and $T \setminus I$ in the other closed half-plane. Hence $\operatorname{conv} A$ lies in the latter closed half-plane. Compactness of $B \subset I$ gives a strictly positive distance of B from that line. The same strict linear inequality holds for every convex combination of points of B. The hulls are disjoint.

Choose an affine linear function ℓ which vanishes on the chord line and is positive on I. Compactness gives $m = \min_{b \in B} \ell(b) > 0$, while $\ell(a) \le 0$ on A. For the respective convex combinations,

$$\ell\left(\sum_j t_j a_j\right) \le 0, \qquad \ell\left(\sum_k s_k b_k\right) \ge m.$$

If the linear part of ℓ has Euclidean norm $M > 0$, these inequalities even give $|u - v| \ge m/M$ for points u, v in the two hulls. The number depends on this pair of compact sets; no uniform separation over all classes of a partition is being asserted.

If an alternating quadruple exists, the straight chords $[a_1, a_2]$ and $[b_1, b_2]$ intersect in $\mathbb{D}$. To see this directly, the chord line through a_1, a_2 separates b_1, b_2, so their segment meets the line. The intersection lies inside the disk, whose intersection with that line is exactly $[a_1, a_2]$. Both segments belong to the respective convex hulls. Thus disjoint hulls rule out alternating quadruples. Finally κ is a bijection, so the last two assertions are equivalent.

For completeness, if the signed distances of b_1, b_2 from the first chord line are $d_1 > 0$ and $d_2 < 0$, then

$$t = \frac{d_1}{d_1 - d_2} \in (0, 1), \qquad w = (1 - t)b_1 + tb_2$$

lies on that line. Since b_1, b_2 are distinct circle points and the combination is strict, $|w| < 1$. Thus the intersection is in the interiors of both chords, rather than merely a possible common endpoint. This is the exact geometric event excluded by the four-distinct-endpoint convention.

□

The criterion is symmetric even though its complementary-arc formulation singles out A. Compactness is used for disjoint *closed* hulls. Without it, hull closures can meet at limiting endpoints absent from the original sets.

12.3 Why conformal boundary fibers are unlinked

Let $K \subset \mathbb{C}$ be a nondegenerate locally connected full continuum. The exterior Riemann map is normalized by

$$F : \widehat{\mathbb{C}} \setminus \overline{\mathbb{D}} \longrightarrow \widehat{\mathbb{C}} \setminus K, \qquad F(\infty) = \infty.$$

The Carathéodory boundary extension established in Volume VIII, combined with the locally connected boundary conclusion for a full locally connected continuum, supplies a continuous extension to T. Write $b = F|_T$. Its fibers are compact and its image is ∂K. This use of the boundary-extension theorem is explicit; it is not a consequence of the elementary quotient lemmas alone.

Theorem 12.6 (Unlinked landing fibers)**.** *Different fibers of b have disjoint closed hyperbolic convex hulls.*

Proof. Suppose instead that fibers over distinct points $x, y \in \partial K$ contain an alternating quadruple a_1, b_1, a_2, b_2 as in Definition 12.4. The two external rays $F(\{ra_j : r > 1\})$ land at x. Including x and ∞, their union is a Jordan curve Γ in the sphere: each ray is embedded, their interiors are disjoint by injectivity of F, and their only common end values are x and ∞.

The rays at b_1, b_2 lie on different sides of Γ near infinity. One can verify the sides on any sufficiently large exterior equipotential $F(\{|z| = R\})$: the four intersection points have the same circular order as their preimages, and the two arcs between the a_j intersections contain the b_j intersections separately. Equivalently, the two radial sectors in the source exterior are the two local sides at the point ∞.

Each b_j ray together with its landing point y is a connected subset of the complement of Γ. It misses Γ in the exterior by injectivity, and it misses $\Gamma \cap K = \{x\}$ because $y \neq x$. Therefore its entire connected image lies in one component of the complement. Both components would contain y, contradicting the Jordan separation theorem. Alternation is impossible.

Writing the two Jordan complementary components as D_1, D_2 makes the contradiction explicit. The initial portion of the ray at b_j lies in D_j. Its entire ray with landing point is connected and avoids Γ, so it is contained in that same component:

$$F(\{rb_j : r > 1\}) \cup \{y\} \subset D_j, \qquad j = 1, 2.$$

This would imply $y \in D_1 \cap D_2 = \varnothing$. The landing point is included here because $y \neq x$; if both fibers had been the same fiber, it would lie on Γ and this argument would correctly cease to apply.

Lemma 12.5 now gives disjoint hulls. □

This proof establishes the precise assertion of source Lemma 9.27. It uses two rays from each of two fibers, so it handles infinite and uncountable fibers without joining their infinitely many rays into an auxiliary continuum. The source's reference to a proposition numbered 2.65 points to an exercise in the inspected edition; the needed separation argument has been supplied here.

12.4 Worked investigations on hull geometry

Investigation 12.7 (A formula for one ideal geodesic)**.** Let $0 < \alpha < \pi/2$ and let the endpoints be $e^{i\alpha}$ and $e^{-i\alpha}$. Compute the geodesic circle and its point on the real axis inside $\mathbb{D}$.

Solution. Symmetry makes its center real, and orthogonality gives radius $\sqrt{c^2-1}$. Substitution of an endpoint gives $2c\cos\alpha = 2$, so

$$c = \sec\alpha, \qquad R = \tan\alpha, \qquad x_{\min} = \sec\alpha - \tan\alpha = \frac{\cos\alpha}{1+\sin\alpha}.$$

The other real intersection is outside the unit disk. In Klein coordinates the geodesic is the vertical chord $\operatorname{Re} w = \cos\alpha$, which also gives $x_{\min} = \kappa^{-1}(\cos\alpha)$.

Investigation 12.8 (A two-point class is not a filled region)**.** Explain why collapsing the hull of $\{-1, 1\}$ is permitted by the disk-decomposition construction although that hull has empty interior.

Solution. Its hull is the diameter $[-1, 1]$, a compact connected set whose complement in the sphere is connected: the two sides can be joined around an endpoint outside the closed disk. The decomposition theorem will require a non-separating continuum, not a set with positive area. The diameter separates the *open disk* but does not separate the *sphere*. Confusing these two ambient spaces would incorrectly exclude exactly the simplest pinching operation.

The ambient distinction can also be seen by explicit paths. In $\mathbb{D} \setminus [-1, 1]$, the sign of the imaginary part cannot change along a path. In $\widehat{\mathbb{C}} \setminus [-1, 1]$, the upper point i and lower point $-i$ can instead be joined along the right semicircle of radius two, using short paths from the endpoints to that semicircle. Those paths pass around 1, while remaining disjoint from the segment. Thus separation inside the source disk is compatible with the nonseparation in the sphere required by Moore.

Investigation 12.9 (Crossing antipodal classes)**.** Why does the closed circle relation $z \sim -z$ fail to give disjoint hyperbolic hulls?

Solution. Every class has a diameter as its hull. Distinct diameters meet at 0, so these hulls cannot be different classes of a partition of $\overline{\mathbb{D}}$. For example, $\{1, -1\}$ and $\{i, -i\}$ alternate. The relation is closed and its circle quotient is a circle, but the unlinked hypothesis needed for the disk filling fails.

Investigation 12.10 (Closedness of a fiber relation)**.** Suppose $b : T \to Z$ is continuous and Z is Hausdorff. Show directly that its fiber relation is closed and explain why a discontinuous choice of a representative from each fiber is irrelevant.

Solution. If $s_n \to s$, $t_n \to t$ and $b(s_n) = b(t_n)$, continuity gives $b(s) = b(t)$. This is the sequential closedness criterion on T^2. No representative-selection map occurs. The quotient is specified by all pairs in each fiber; choosing a preferred representative is an additional problem and need not have a continuous solution.

Investigation 12.11 (One-sided limits of hulls)**.** Let A_n be classes of a fixed closed relation on T, and suppose $a_n \in A_n$ tends to a. Prove that every limit of points $z_n \in H(A_n)$ lies in $H([a])$.

Solution. Write

$$\kappa(z_n) = \sum_{j=1}^{3} t_{n,j} a_{n,j}, \qquad a_{n,j} \in A_n.$$

Pass to a subsequence along which the three coefficients and three circle points converge. Every pair $(a_n, a_{n,j})$ belongs to the closed relation, so each limit a_j belongs to $[a]$. Passing to the limit gives $\kappa(z) = \sum_j t_j a_j \in \operatorname{conv}[a]$. The claim follows by the definition of H. This argument proves an upper-limit inclusion, not equality with the whole limiting hull.

Chapter 13

A precise pinched-disk construction

Source correspondence. Lyubich, The closed-decomposition content of Theorem 9.28.

We now start from an *actual closed equivalence relation* R on T, whose different classes are unlinked. No dynamics is assumed. The hypotheses are strong enough to construct a compact planar model without making an unspecified choice of a cleaned lamination.

For each class A let $H(A)$ be its closed ideal hull. The preceding chapter proves that these hulls are compact, connected and pairwise disjoint. Declare each $H(A)$ to be one class on $\overline{\mathbb{D}}$ and leave every point outside their union a singleton. Extend this partition to $\widehat{\mathbb{C}}$ by singleton classes outside $\overline{\mathbb{D}}$. We denote the resulting relation by $\widehat{R}$ and put

$$S = \bigcup_{A \in T/R} H(A), \qquad G = \mathbb{D} \setminus S.$$

The open components of G will be called gaps. The set S includes the whole circle, including singleton boundary classes.

For the finite illustration below, one may take the three-point class at angles $0, 2\pi/3, 4\pi/3$, a two-point class at $\pi/6, \pi/3$, and singleton classes everywhere else. The latter chord lies wholly in one cap of the triangle, so the two classes are unlinked. The relation is closed because its graph is the diagonal together with finitely many additional pairs. In particular,

$$\widehat{R} = \Delta_{\widehat{\mathbb{C}}} \cup \bigcup_{A \in T/R} \big(H(A) \times H(A)\big).$$

This formula describes the fibers, but the union may have infinitely many terms in general. Its closedness must still be proved; arbitrary unions of closed sets need not be closed.

13.1 Closedness of the filled relation

Theorem 13.1 (Closed hull decomposition)**.** *The set S is closed in $\overline{\mathbb{D}}$, and $\widehat{R}$ is a closed equivalence relation on the sphere. Its classes are connected and do not separate the sphere. On T, the restriction of $\widehat{R}$ is exactly R.*

Proof. Disjointness of the hulls makes the proposed classes a partition. The equality $H(A) \cap T = A$ follows from Lemma 12.2 and the fact that κ fixes T. Thus the restriction to T is the specified relation.

The partition assertion includes points not contained in any hull. Such a point is assigned its own singleton, so every point of the sphere lies in at least one class. If it lay in two, either it would be in two different hulls, contrary to the chord criterion, or its singleton would have

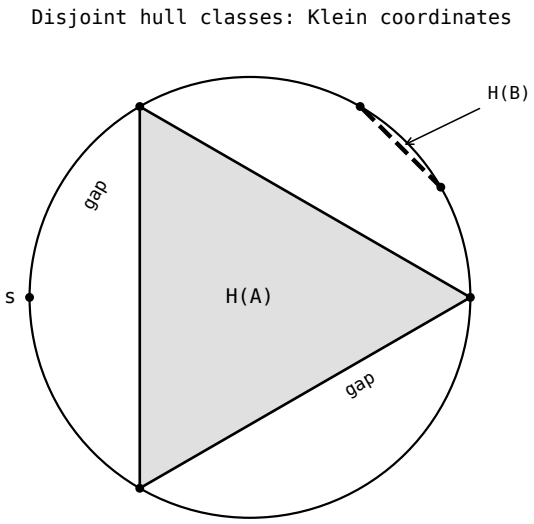


Figure 13.1: The disk in Klein coordinates: a filled three-point class, a disjoint two-point chord class, and a marked singleton. All other unmarked circle points are singletons. The complementary open regions are gaps. Straight chords in this coordinate represent hyperbolic geodesics in the original disk.

been assigned outside the prescribed remaining set. Thus the construction has unique classes. In particular, transitivity is inherited from a partition:

$$z\widehat{R}w, \quad w\widehat{R}v \quad \Longrightarrow \quad z, w, v \text{ lie in the unique class containing } w.$$

There is no transitive-closure operation hidden here.

To prove that S is closed, suppose $z_n \in H(A_n)$ and $z_n \to z$. Choose $a_n \in A_n$ and pass to a subsequence with $a_n \to a \in T$. The last investigation in the preceding chapter, or its three-point argument, gives

$$z \in H([a]_R) \subset S.$$

Every convergent sequence in S therefore has its limit in S.

The anchor a_n is used to keep all limiting vertices in one class. Explicitly, choose

$$\kappa(z_n) = \sum_{j=1}^{3} t_{n,j} u_{n,j}, \quad u_{n,j} \in A_n, \quad (t_{n,1}, t_{n,2}, t_{n,3}) \in \Delta_2.$$

Compactness of $T^4 \times \Delta_2$ gives a single subsequence on which the anchor, all three vertices and all coefficients converge. For each j, closedness gives aRu_j from $a_nRu_{n,j}$. The barycenter limit then lies in $\mathrm{conv}[a]_R$. It is enough to extract such a subsequence because the original sequence already has the unique prescribed limit z.

For closedness of $\widehat{R}$, let $z_n\widehat{R}w_n$ and $(z_n, w_n) \to (z, w)$. If infinitely many pairs lie on the diagonal, then $z = w$, which is sufficient. Otherwise pass to a subsequence for which z_n, w_n belong to the same hull $H(A_n)$. Choose $a_n \in A_n$ and extract a subsequence converging to a. Apply the three-point argument simultaneously to z_n and w_n:

$$\kappa(z_n) = \sum_{j=1}^{3} t_{n,j} a_{n,j}, \qquad \kappa(w_n) = \sum_{j=1}^{3} s_{n,j} b_{n,j}.$$

Compactness permits all six circle points and both coefficient vectors to converge after passing to one subsequence. Closedness of R and $a_nRa_{n,j}$, $a_nRb_{n,j}$ put all six limits in $[a]_R$. Hence

$z, w \in H([a]_R)$, as required. Points outside $\overline{\mathbb{D}}$ create no further case, because a sequence of non-singleton classes stays inside that compact disk.

Here the full simultaneous parameter space is

$$T^7 \times \Delta_2 \times \Delta_2 : \quad (a_n, u_{n,1}, u_{n,2}, u_{n,3}, v_{n,1}, v_{n,2}, v_{n,3}; t_n, s_n).$$

One anchor and six vertices suffice. Separate subsequences for the two barycenters would not by themselves guarantee a common limiting class. There is also a useful case check: if $z \neq w$, their spherical distance is positive, so for all sufficiently large n,

$$d_{\mathrm{sph}}(z_n, w_n) > \tfrac{1}{2} d_{\mathrm{sph}}(z, w) > 0.$$

Those pairs must be in hull classes rather than singleton classes. Both limits then lie in $\overline{\mathbb{D}}$. If a limit is outside that disk, or is infinity, only the diagonal case can occur.

Every Euclidean compact convex set C has connected complement in the sphere. Here is a direct verification. Choose $c \in C$. For a point $z \notin C$, the ray $c + t(z - c)$, $t \geq 1$, remains outside C: if a later point were in C, convexity would place z between two points of C. Follow this ray to a circle centered at c large enough to enclose C. Any two outside points can be connected using their outward rays and an arc of such a circle. Thus the complement is path connected. The assertion also holds when C is a singleton.

The outward-ray claim uses convexity in the following exact form. If $y = c + t(z - c) \in C$ for $t > 1$, then

$$z = (1 - t^{-1})c + t^{-1} y \in C,$$

contrary to the choice of z. For two points outside C, choose a radius R about c greater than both of their distances from c and greater than $\max_{u \in C} |u - c|$. Each outward segment reaches that circle without entering C, and the whole circle misses C. Their concatenation is the asserted path. This also handles segments and singleton convex sets.

Extend κ to the sphere by the identity outside $\mathbb{D}$. The two definitions agree on T, so this is a sphere homeomorphism. It carries each $H(A)$ to $\operatorname{conv} A$. Thus the complement of each hull is connected. Singleton classes outside the disk are also non-separating. Connectedness of the classes follows from convexity and continuity of κ^{-1}. □

The proof uses compactness of *one fixed closed relation.* A sequence of unrelated finite partitions need not have an equivalence relation as its limit; the final chapter supplies a concrete example.

Proposition 13.2 (Continuity of the collapsed boundary map)**.** *Let $b : T \to Z$ be continuous into a Hausdorff space and suppose its fibers are exactly the classes of R. The assignment*

$$\beta(z) = b(a) \quad \text{if } z \in H([a]_R)$$

defines a continuous map $\beta : S \to b(T)$.

Proof. The assignment is well defined by disjointness of the hulls and constancy of b on each boundary class. The target $b(T)$ is compact and Hausdorff. It suffices to prove the closed-graph property for β and then use compactness: if the graph is closed, the inverse image of a closed target set is the projection of a compact part of the graph and hence closed.

Let $z_n \to z$ in S and suppose $\beta(z_n) \to y$ in $b(T)$. Choose a_n in the boundary class defining z_n and pass to a subsequence with $a_n \to a$. Continuity gives $b(a) = y$. The hull limit argument gives $z \in H([a]_R)$, and therefore $\beta(z) = b(a) = y$. Since $S \times b(T)$ is first countable in this case, the sequential argument proves the graph closed. Alternatively, $b(T)$ is homeomorphic to the compact quotient T/R and has a countable base by Proposition 11.4; a product of two first-countable spaces is first countable. Thus the sequential criterion used here needs no unmentioned arbitrary-space assumption.

To finish the closed-graph step in set notation, let $\Gamma_\beta \subset S \times b(T)$ be the graph and let $F \subset b(T)$ be closed. Then

$$\beta^{-1}(F) = \pi_S(\Gamma_\beta \cap (S \times F)).$$

The set inside the projection is compact, since the graph is closed in the compact product. Its projection is compact and hence closed in S. Thus every inverse image of a closed target set is closed. This is why a closed graph implies continuity here: the compact target and compact graph supply a closed projection, a property which cannot be assumed for an arbitrary graph.

□

13.2 The external sphere theorem

The remaining planar recognition step is a classical topological input, not a consequence of convexity alone. We use the following form of Moore's theorem: a nontrivial closed decomposition of the two-sphere into connected non-separating compact sets has a quotient homeomorphic to the two-sphere. Here nontrivial means that the entire sphere is not a single class. Timorin's exposition states this as Theorem 1.1 and develops a proof [2]. We use the theorem as a named external input; we do not count its proof as a result established in this volume.

The convex-hull extension has also been used in conformal-boundary interpolation; Bishop's Lemma 13 states an upper-semicontinuity extension of this kind [3]. Our Theorem 13.1 supplies the compactness argument needed here directly, including the two-point and singleton cases.

Theorem 13.3 (The planar model of a closed unlinked relation)**.** *For a closed unlinked relation R on T, there is a continuous surjection $p : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ with precisely the $\widehat{R}$-classes as fibers. It can be chosen with $p(\infty) = \infty$. The set*

$$L = p(\overline{\mathbb{D}})$$

is a locally connected full continuum, possibly a singleton. The following identities and restrictions hold:

$$p^{-1}(L) = \overline{\mathbb{D}}, \qquad \partial L = p(T) = p(S), \qquad \operatorname{int} L = p(G).$$

The restriction of p to the exterior of $\overline{\mathbb{D}}$ is a homeomorphism onto the exterior of L, and its restriction to G is a homeomorphism onto $\operatorname{int} L$.

Proof. Let $q : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}/\widehat{R}$ be the quotient map. Theorem 13.1 verifies each hypothesis of Moore's theorem. The relation is not universal because, for example, 2 and 3 are different singleton classes. Let h be the resulting homeomorphism from the quotient onto the sphere. Postcompose with a sphere homeomorphism sending $h(q(\infty))$ to ∞, and set $p = h \circ q$.

The hypothesis check separates four facts:

$$\widehat{R} \text{ is closed}, \qquad [z]_{\widehat{R}} \text{ is compact and connected},$$
$$\widehat{\mathbb{C}} \setminus [z]_{\widehat{R}} \text{ is connected}, \qquad [2]_{\widehat{R}} \neq [3]_{\widehat{R}}.$$

The first comes from the simultaneous limiting argument, the next two from convexity and the sphere coordinate, and the last from the singleton exterior convention. Only after these checks does the named Moore theorem identify the quotient as a sphere. It does not assert that the quotient map is conformal or supply a numerical modulus for the homeomorphism h.

The disk is saturated, so $p^{-1}(L) = \overline{\mathbb{D}}$. Thus L is compact and misses ∞; it is a planar continuum because the closed disk is connected. Its complement is the image of the connected exterior and hence is connected. Local connectivity follows from Theorem 11.7 applied to the restriction $p : \overline{\mathbb{D}} \to L$.

Every point outside S has a singleton fiber. The exterior and G are open saturated sets, so p is a homeomorphism on each of them onto its image. To verify the inverse-continuity assertion, an open subset of either set is also open in the sphere and is saturated; its image is open in the quotient. In particular, $p(G)$ is an open subset of L.

More explicitly, if O is open in G, then it is open in the sphere because G itself is open. Every fiber meeting O is its one point, so

$$p^{-1}(p(O)) = O.$$

The quotient definition makes $p(O)$ open in the sphere, not merely relatively open in L. Thus $p|_G$ is a continuous bijection which is open, and its inverse is continuous. The same calculation applies on the exterior. An arbitrary open subset meeting a collapsed hull would not have this saturation identity.

For $\zeta \in T$, choose $r_n > 1$ with $r_n \downarrow 1$. Then

$$p(r_n\zeta) \notin L, \qquad p(r_n\zeta) \longrightarrow p(\zeta) \in L.$$

Hence $p(T) \subset \partial L$. Conversely, every point of L comes from S or G, and every hull in S meets T. Thus

$$L = p(S) \cup p(G) = p(T) \cup p(G),$$

where the first set lies in the boundary and the second lies in the interior. The two sets are disjoint because a singleton in G is not a hull class. This proves all the stated identities.

For example, the reverse boundary inclusion follows by an exact exclusion:

$$y \in \partial L \implies y \in L \setminus p(G) \implies y \in p(T).$$

The first implication uses that L is closed and $p(G)$ is open in the sphere. Conversely, every circle image is approached by exterior images, so cannot be an interior point of L. Combining these facts also gives $L \setminus \partial L = p(G)$. No inference about the boundary of a continuous image is being made without these two separate inclusions.

□

The exterior homeomorphism obtained here is topological. It is not asserted to be conformal, quasiconformal, Lipschitz or effectively computable. Identifying this model with a specified conformal hull requires an additional matching argument, given in the next chapter.

Corollary 13.4 (Boundary and gap interpretation)**.** *The boundary ∂L is homeomorphic to T/R. The components of G correspond bijectively, under p, to the components of* $\operatorname{int} L$*. In particular there are at most countably many gaps.*

Proof. The boundary assertion is Proposition 11.3, applied to $p|_T$. The component assertion follows from the homeomorphism $G \to \operatorname{int} L$. Each open component of $G \subset \mathbb{R}^2$ contains a point with rational coordinates, and distinct components contain different such points after making one choice per component. Thus there are at most countably many. □

13.3 Worked investigations on the disk model

Investigation 13.5 (The universal boundary class)**.** Take all of T as one class. Determine S, G and the model L. Does this violate Moore's nontriviality hypothesis?

Solution. The convex hull of T is $\overline{\mathbb{D}}$, so $S = \overline{\mathbb{D}}$ and $G = \varnothing$. The entire disk is collapsed and L is one point. All exterior points remain separate. In particular the sphere is not one class, so Moore's hypothesis holds. The quotient sphere is still a sphere; the subset singled out as its pinched-disk model has become degenerate.

Investigation 13.6 (One pinched diameter)**.** Take $\{-1, 1\}$ as the only non-singleton boundary class. Describe the gaps and the topology of L.

Solution. Here $S = T \cup [-1, 1]$, and G has two components, the open upper and lower half-disks. Their images are the two interior components of L. Collapsing the straight edge of a closed half-disk gives a closed disk: map the half-disk first to a triangle with that edge as base, then identify the base to the cone vertex; equivalently use a homeomorphism to a square that takes the chosen boundary arc to one edge and use the quotient of $[0, 1]^2$ collapsing $[0, 1] \times \{0\}$, which is a triangle. Explicitly, $(s, t) \mapsto (st, t)$ maps the square onto $\{(u, t) : 0 \le u \le t \le 1\}$ and its only non-singleton fiber is the bottom edge. Compact quotient factorization therefore identifies that quotient with the triangle. The two resulting disks meet precisely at the image of the diameter. Thus L is a union of two closed disks touching at one point. Their interiors are distinct and their boundary contact is a single pinch, not a boundary arc.

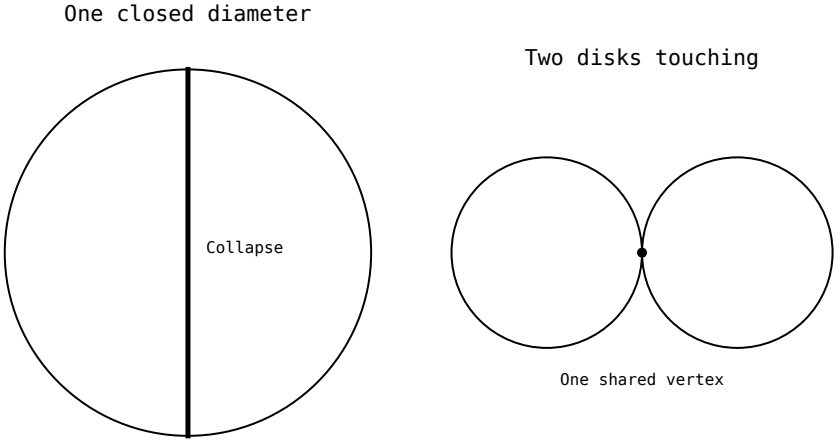


Figure 13.2: Collapsing one closed diameter produces two closed disks meeting at one point. The right-hand drawing specifies the quotient's topology; it does not preserve Euclidean lengths or a conformal metric.

Investigation 13.7 (One triangular class)**.** Let A consist of three distinct circle points and leave all other boundary points singleton. How many interior components does the model have, and what is their common intersection?

Solution. In Klein coordinates the hull is the filled inscribed triangle. Its complement in the open disk has three cap components, one for each complementary circle arc. Each closed cap has one chord edge collapsed, so the preceding square-to-triangle argument makes its quotient a closed disk. The three disks meet at the one point $p(H(A))$, and no other intersection is possible because different cap closures meet only within the triangular hull. Thus the model has three interior components and a common pinched point. A filled triangular class is collapsed in its entirety; collapsing only its three sides would leave an unintended interior.

Investigation 13.8 (An interval from a filled disk)**.** Take the circle relation whose classes are $\{e^{it}, e^{-it}\}$. Give an explicit quotient coordinate on the closed disk and identify the model.

Solution. In Klein coordinates the classes fill the disk by the vertical segments with constant real part. Thus

$$u(z) = \operatorname{Re} \kappa(z) = \frac{2 \operatorname{Re} z}{1 + |z|^2}$$

has precisely the filled classes as fibers, and maps the disk onto $[-1, 1]$. Proposition 11.3 identifies the disk quotient with this interval. Every point of the disk lies in a filled class, so $G = \varnothing$ and the model has empty interior. A quotient of a two-dimensional disk need not itself have two-dimensional interior.

Investigation 13.9 (Non-small classes are allowed)**.** Show that the non-singleton classes of a valid closed hull decomposition need not form a null family.

Solution. Use the relation in Investigation 13.8. In Klein coordinates the class with real coordinate x is the vertical segment of length $2\sqrt{1 - x^2}$. For $|x| \leq 1/2$ these lengths are at least $\sqrt{3}$. Their inverse images under κ cannot have diameters tending to zero, since uniform continuity of κ would then force their image diameters to tend to zero. There are uncountably many such disjoint classes. Upper semicontinuity, not a null-family condition on the classes, is the correct decomposition hypothesis. Null families of complementary *domains*, used in the recognition theorem, are a different matter.

Chapter 14

Recognizing a locally connected conformal hull

14.1 The boundary correspondence determines the outside

Let $K \subset \mathbb{C}$ be a nondegenerate locally connected hull. Fix its normalized exterior conformal map

$$F : \widehat{\mathbb{C}} \setminus \overline{\mathbb{D}} \longrightarrow \widehat{\mathbb{C}} \setminus K, \qquad F(\infty) = \infty.$$

The boundary-coordinate input, Theorem 2.5, gives a continuous extension to $\widehat{\mathbb{C}} \setminus \mathbb{D}$, with boundary map $b : T \to \partial K$. Its fiber relation is closed and unlinked by Theorem 12.6. Apply the construction of Theorem 13.3 to this relation. Write

$$p : \widehat{\mathbb{C}} \longrightarrow \widehat{\mathbb{C}}, \qquad L = p(\overline{\mathbb{D}}), \qquad p(\infty) = \infty,$$

for the resulting Moore projection and model hull. The choices in identifying the quotient sphere with the standard sphere do not affect the equivalence relation or its marked boundary.

The remaining task is recognition: why does this abstract hull model the given K, including its interior? Merely knowing $\partial L \cong \partial K$ is not yet a proof that a prescribed boundary homeomorphism extends across every complementary region. We establish the extension and its continuity at accumulating regions explicitly.

Lemma 14.1 (The exterior correspondence)**.** *There is a homeomorphism*

$$H_0 : \widehat{\mathbb{C}} \setminus \operatorname{int} L \longrightarrow \widehat{\mathbb{C}} \setminus \operatorname{int} K \quad \textit{such that} \quad H_0(p(z)) = F(z) \quad (|z| \geq 1).$$

In this formula the point ∞ is included. Its boundary restriction $h : \partial L \to \partial K$ satisfies $h(p(\zeta)) = b(\zeta)$.

Proof. Put $A = \widehat{\mathbb{C}} \setminus \mathbb{D}$, a compact set. From Theorem 13.3,

$$p(A) = (\widehat{\mathbb{C}} \setminus L) \cup \partial L = \widehat{\mathbb{C}} \setminus \operatorname{int} L.$$

The restriction $p|_A$ is a continuous surjection between compact Hausdorff spaces and therefore a quotient map. The extension of F maps A onto $\widehat{\mathbb{C}} \setminus \operatorname{int} K$: the open exterior maps bijectively onto $\widehat{\mathbb{C}} \setminus K$, while the circle maps onto ∂K. An exterior image cannot be a boundary image. On the open exterior p and F both have singleton fibers. On T both have exactly the classes of the relation defined by b. Consequently

$$p(z) = p(w) \quad \Longleftrightarrow \quad F(z) = F(w) \qquad (z, w \in A).$$

The factorization theorem, Proposition 11.3, supplies a continuous bijection H_0 with the asserted equation. Its domain is compact and its target Hausdorff, so it is a homeomorphism.

There are exactly three fiber cases on the closed exterior. Two open-exterior points have equal images only if they are equal, by conformal injectivity. One exterior and one circle point cannot have equal images because their images lie respectively in $\widehat{\mathbb{C}} \setminus K$ and ∂K. Two circle points have equal images exactly when they belong to the specified boundary class. Hence the map can be written without a choice of inverse branch as

$$F\big(p^{-1}(p(z)) \cap A\big) = \{F(z)\}, \qquad H_0(p(z)) = F(z).$$

The class notation is the content of factorization: every representative gives the same right-hand side.

□

This argument includes a boundary point having several external accesses. One never chooses an inverse branch of b there. The equation defining H_0 is well defined because all choices are exactly the fibers which have already been collapsed.

14.2 Matching and filling the Jordan regions

Lemma 14.2 (Matching interior components)**.** *The boundary homeomorphism h induces a bijection between the components U of* $\operatorname{int} L$ *and the components V of* $\operatorname{int} K$*, characterized by*

$$h(\partial U) = \partial V.$$

For each matched pair, $h|_{\partial U}$ extends to a homeomorphism $H_U : \overline{U} \to \overline{V}$.

Proof. Both L and K are locally connected hulls. Theorem 8.3 makes every interior component a Jordan domain. Fix one such U. The curve $\gamma = h(\partial U)$ is therefore a Jordan curve in ∂K. Its bounded complementary domain V lies in K by Lemma 2.2. Since V is open, it lies in $\operatorname{int} K$.

It is a whole component of $\operatorname{int} K$. Indeed, a connected subset of $\operatorname{int} K$ meeting both the inside and the outside of γ would meet γ, by Jordan separation. This is impossible because $\gamma \subset \partial K$. The connected open set V is thus precisely the component which it meets. Applying the same argument to h^{-1} shows that every component of $\operatorname{int} K$ arises in this way. It also proves injectivity: the Jordan curve ∂V has only one bounded complementary domain, so applying the inverse construction recovers U.

In detail, let C be the interior component containing one point of V. Because V is connected and lies in $\operatorname{int} K$, it lies in C. On the other hand, C misses γ and is connected, so Jordan separation forces C to stay on the bounded side:

$$V \subset C \subset V.$$

Thus $C = V$, not merely an open subdomain of that component. For the inverse construction, $h^{-1}(\partial V) = \partial U$, and the bounded side of this Jordan curve is exactly U. This proves both surjectivity and injectivity of the component matching.

Choose closed-disk parametrizations

$$a : \overline{\mathbb{D}} \longrightarrow \overline{U}, \qquad c : \overline{\mathbb{D}} \longrightarrow \overline{V}$$

which are homeomorphisms. They exist by the Jordan-domain theorem. The circle map $\eta = c^{-1} \circ h \circ a|_T$ is a homeomorphism of T. Its radial extension is

$$E_\eta(0) = 0, \qquad E_\eta(r\zeta) = r\eta(\zeta) \quad (0 < r \le 1,\ \zeta \in T).$$

This map is continuous away from zero and continuous at zero because $|E_\eta(z)| = |z|$. Its inverse is $E_{\eta^{-1}}$. Consequently

$$H_U = c \circ E_\eta \circ a^{-1}$$

is the required homeomorphism, and the formula on the boundary reduces exactly to h.

A useful continuity estimate for the radial extension is

$$|r\eta(\zeta) - s\eta(\xi)| \le |r - s| + \min\{r, s\}\, |\eta(\zeta) - \eta(\xi)|.$$

For $r \ge s$, subtract and add $s\eta(\zeta)$; the other case is symmetric. Away from the center, angular coordinates depend continuously on the point, so the estimate proves continuity there. At the center it reduces to $|E_\eta(r\zeta)| = r$ uniformly in ζ. The identities

$$E_{\eta^{-1}} E_\eta = \mathrm{id}_{\overline{\mathbb{D}}}, \qquad c E_\eta a^{-1}|_{\partial U} = h|_{\partial U}$$

then verify both the inverse and the boundary agreement.

□

No conformality is asserted for H_U. A boundary homeomorphism need not be the boundary value of a conformal equivalence between the two domains. The radial extension solves the topological problem required for recognition; it does not introduce an unproved analytic interpolation condition.

14.3 Continuity across infinitely many regions

Theorem 14.3 (Recognition of the locally connected model)**.** *Let K and F be as above, and let $\widehat{R}$ be the filled-hull equivalence relation obtained from the fibers of $F|_T$. There is a homeomorphism of pairs*

$$(\widehat{\mathbb{C}}/\widehat{R},\ \overline{\mathbb{D}}/\widehat{R}) \cong (\widehat{\mathbb{C}}, K).$$

It can be chosen so that the associated continuous surjection $\Phi : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ satisfies

$$\Phi(z) = F(z) \quad (|z| \ge 1), \qquad \Phi(\overline{\mathbb{D}}) = K,$$

and has exactly the filled sets $H(A)$ and the remaining singleton points as fibers.

Proof. Use Lemma 14.1 on $A_0 = \widehat{\mathbb{C}} \setminus \operatorname{int} L$. For every interior component U of L, choose an extension H_U from Lemma 14.2. Define

$$H(x) = \begin{cases} H_0(x), & x \in A_0, \\ H_U(x), & x \in U. \end{cases}$$

These formulas agree wherever their closed domains overlap. Distinct component closures can meet only on ∂L, and all boundary values are the same map h. The bijection of components and the bijectivity of each restriction imply that H is a bijection from the sphere to itself.

One can verify global injectivity by the disjoint decompositions

$$\widehat{\mathbb{C}} = A_0 \sqcup \bigsqcup_j U_j, \qquad \widehat{\mathbb{C}} = (\widehat{\mathbb{C}} \setminus \operatorname{int} K) \sqcup \bigsqcup_j V_j.$$

The first part is mapped bijectively onto the first target part by H_0; each remaining part is mapped bijectively onto its matched part. Images of different parts cannot collide. At the level of closures, intersections lie on the common boundary where every formula equals h. Thus

compatibility on overlaps and bijectivity of the assembled map are separate checks, both provided by the construction.

It remains to prove continuity. There are at most countably many components, because each nonempty open component contains a point of a fixed countable dense subset of the plane. If the components are finite in number, the finite closed-set pasting argument applies immediately. In general enumerate them as U_j, with matched target components V_j. The null-family theorem 3.5 applies separately to L and K:

$$\operatorname{diam} \overline{U}_j \longrightarrow 0, \qquad \operatorname{diam} \overline{V}_j \longrightarrow 0 \quad \text{along distinct indices } j.$$

All these sets lie in bounded subsets of the plane, so either Euclidean or spherical diameters give the required conclusion.

The null statement is independent of the enumeration. For every $\varepsilon > 0$ the two exceptional sets

$$\{j : \operatorname{diam} \overline{U}_j \geq \varepsilon\}, \qquad \{j : \operatorname{diam} \overline{V}_j \geq \varepsilon\}$$

are finite. The second set is finite even though the matching might permute the target components in a complicated order: a bijection takes a finite set of large target components back to a finite set of indices. Their union is still finite. Outside it, both the source boundary approximation and the target image diameter are simultaneously smaller than ε.

For every $x \in U_j$,

$$\operatorname{dist}(x, \partial U_j) \leq \operatorname{diam} \overline{U}_j, \qquad \operatorname{diam} H(\overline{U}_j) = \operatorname{diam} \overline{V}_j.$$

These are precisely the two smallness hypotheses of Lemma 3.6, with closed set A_0.

For clarity, the only new sequential case can also be written directly. Suppose $x_n \in U_{j_n}$, the indices are distinct, and $x_n \to x$. Choose $a_n \in \partial U_{j_n}$ with $|x_n - a_n| \leq \operatorname{diam} \overline{U}_{j_n}$. Then $a_n \to x$, and the closed set ∂L contains x. Therefore

$$|H(x_n) - h(x)| \leq \operatorname{diam} \overline{V}_{j_n} + |h(a_n) - h(x)| \longrightarrow 0.$$

Sequences in one fixed closed component are controlled by its chosen extension; sequences in A_0 are controlled by H_0. These alternatives establish continuity everywhere, including accumulation points of infinitely many different components.

The sequential estimate has a direct epsilon interpretation. Let ω_h be the modulus of uniform continuity of h on the compact boundary. For $u \in U_j$ and $x \in \partial L$, choosing $a \in \partial U_j$ yields

$$|H(u) - h(x)| \leq \operatorname{diam} \overline{V}_j + \omega_h(|u - x| + \operatorname{diam} \overline{U}_j).$$

Given $\varepsilon > 0$, choose $\delta > 0$ such that $\omega_h(2\delta) < \varepsilon/2$. Outside finitely many components, require $\operatorname{diam} \overline{U}_j < \delta$ and $\operatorname{diam} \overline{V}_j < \varepsilon/2$. For these components, $|u - x| < \delta$ gives $|H(u) - h(x)| < \varepsilon$. The finitely many exceptional components are handled by their individual continuous extensions. No common modulus for all of those arbitrarily chosen extensions is assumed.

A continuous bijection of the compact sphere onto the Hausdorff sphere is a homeomorphism. Set $\Phi = H \circ p$. Since H is injective, the fibers of Φ are exactly those of p, namely the prescribed filled hulls and singleton points. The equation on the closed exterior follows from the defining equation of H_0. Finally $H(L) = K$, so $\Phi(\overline{\mathbb{D}}) = K$. Factoring through the quotient gives the claimed homeomorphism of pairs.

The fiber conclusion can be checked in both directions:

$$\Phi(z) = \Phi(w) \iff H(p(z)) = H(p(w)) \iff p(z) = p(w) \iff z \widehat{R} w.$$

Thus the interior fillings neither add identifications between previously distinct classes nor split one previously collapsed class. This exact equivalence is stronger than merely showing that Φ is onto K from the disk. It is what identifies the quotient pair with the desired physical pair.

□

This proves the locally connected pinched-disk assertion of source Theorem 9.28 with its dependencies visible. Moore's sphere-decomposition theorem is a named external input. The boundary extension is the precise Caratheodory input already developed in Volume VIII. The interior Jordan theorem and the null-family estimate have proofs in this volume. In particular, the shrinking of the target interior components is supplied by Theorem 3.5; it is not a consequence of the unlinked-fiber lemma. This replaces the source's mismatched reference in its continuity step.

Corollary 14.4 (The model with no interior)**.** *If* $\operatorname{int} K = \varnothing$*, then the union of the filled classes is all of* $\overline{\mathbb{D}}$*. The map* $\Phi|_{\overline{\mathbb{D}}}$ *is determined by the collapsed boundary values and has no open region on which an additional filling must be chosen.*

Proof. The component bijection gives $\operatorname{int} L = \varnothing$. By Theorem 13.3, p maps $G = \mathbb{D} \setminus \bigcup_A H(A)$ homeomorphically onto $\operatorname{int} L$, so G is empty. The circle is already contained in the union of the hulls. Every point of the closed disk therefore belongs to one of them, and its image is the boundary value shared by that class. □

At the other extreme, if K is a closed Jordan disk, its boundary map is one-to-one. The circle classes are singletons, their union is just T, and $G = \mathbb{D}$. There is a single interior region to fill. The theorem covers both these extremes and arbitrary countable families of intermediate Jordan regions.

14.4 Changing the angular marking

Proposition 14.5 (Topological invariance of the marked model)**.** *Suppose two closed unlinked circle relations* R_1, R_2 *are conjugate by a circle homeomorphism* η*:*

$$\zeta R_1 \xi \quad \Longleftrightarrow \quad \eta(\zeta) R_2 \eta(\xi).$$

Then their model hulls are ambiently homeomorphic, with boundary correspondence induced by η*.*

Proof. Let p_i be their Moore projections and $L_i = p_i(\overline{\mathbb{D}})$. Extend η to the closed exterior by $E_\eta(r\zeta) = r\eta(\zeta)$ for $r \geq 1$ and $E_\eta(\infty) = \infty$. This is a homeomorphism; continuity at infinity follows because it preserves the modulus. The maps p_1 and $p_2 E_\eta$ have exactly the same fibers on that closed exterior. Compact quotient factorization therefore gives a homeomorphism between $\widehat{\mathbb{C}} \setminus \operatorname{int} L_1$ and $\widehat{\mathbb{C}} \setminus \operatorname{int} L_2$.

Both hulls are locally connected and full by Theorem 13.3. The argument of Lemma 14.2, using only these properties and the boundary homeomorphism, matches their interior components. Extend over each matched pair radially in Jordan coordinates. The two null-family estimates and Lemma 3.6 prove global continuity exactly as in the recognition theorem. The resulting bijection of the sphere is a homeomorphism and has the required boundary correspondence. □

Rotations commute with the Klein coordinate and carry filled hyperbolic hulls to filled hyperbolic hulls. An arbitrary circle homeomorphism need not do that by its radial extension inside the disk. Proposition 14.5 asserts equivalence of the quotient models, not a pointwise formula carrying every geodesic by radial transport, and not a metric distortion bound.

14.5 Solved recognition investigations

Investigation 14.6 (What the marking does not measure)**.** Does equality of the angular fiber relations determine the area or logarithmic capacity of a hull?

Solution. No. Every round closed disk has a one-to-one exterior boundary parametrization, hence the trivial angular equivalence relation. Disks of radii $r_1 \neq r_2$ have areas $\pi r_1^2, \pi r_2^2$. Their normalized exterior maps are $F_j(w) = a_j + r_j w$, so their logarithmic capacities, read as the positive leading coefficient of the exterior normalization, are r_j. The relation records identifications and the recognition theorem yields topological equivalence. The discarded conformal scale is exactly what these quantities require.

Investigation 14.7 (A radial extension in angle coordinates)**.** Let an orientation-preserving circle homeomorphism have a strictly increasing continuous lift $\psi : \mathbb{R} \to \mathbb{R}$ satisfying $\psi(t + 2\pi) = \psi(t) + 2\pi$. Verify the extension used above.

Solution. Define

$$E(0) = 0, \qquad E(re^{it}) = re^{i\psi(t)} \quad (0 < r \leq 1).$$

Changing t by $2\pi k$ does not change the value, by the lift identity. Away from zero polar coordinates prove continuity locally; at zero, $|E(z)| = |z|$ proves it directly. The inverse uses the increasing lift ψ^{-1}, so $E^{-1}(re^{is}) = re^{i\psi^{-1}(s)}$. Thus E is a homeomorphism of the disk and restricts to the given circle map. No differentiability of ψ is required.

Investigation 14.8 (Why target diameters matter)**.** Construct continuous functions on shrinking disjoint intervals which agree at all endpoints with zero but whose pasted function is discontinuous at their accumulation point.

Solution. Set

$$I_n = [2^{-n},\, 2^{-n} + 2^{-n-2}] \subset [0, 1] \quad (n \geq 1).$$

The intervals are disjoint and their diameters tend to zero. On each I_n take the piecewise linear tent with value zero at both endpoints and value one at the midpoint m_n; set the function to zero off their interiors. Each closed-interval restriction is continuous and the endpoint data agree with the zero function on the remaining closed set. But $m_n \to 0$ while $f(m_n) = 1$ and $f(0) = 0$. The source diameters vanish and the target diameters remain one. This violates exactly the target smallness condition needed by Lemma 3.6.

Investigation 14.9 (Recognizing the segment)**.** Relate the interval quotient in Investigation 13.8 to the exterior conformal map of a line segment.

Solution. The map $F(w) = (w + w^{-1})/2$ from Investigation 2.7 has boundary value $F(e^{it}) = \cos t$. Its fibers are the conjugate pairs $\{e^{it}, e^{-it}\}$, including the singleton endpoints at $t = 0, \pi$. Their hyperbolic hulls fill the disk. The explicit quotient coordinate

$$\Phi(z) = \frac{2\,\mathrm{Re}\, z}{1 + |z|^2} \quad (|z| \leq 1)$$

agrees with $\cos t$ on the circle. Together with the prescribed F outside, it gives a continuous sphere projection with disk image $[-1, 1]$. The disk map is not injective; its complete fibers are the filled geodesics. The quotient, rather than the uncollapsed disk, is homeomorphic to the segment.

Investigation 14.10 (The freedom inside one gap)**.** Exhibit distinct ambient extensions in the recognition theorem which agree on the closed exterior and on every boundary value.

Solution. Suppose there is a matched component V. In disk coordinates choose a continuous function $\chi : [0, 1] \to \mathbb{R}$ with $\chi(1) = 0$ and with some value not an integer multiple of 2π. The twist

$$T_\chi(0) = 0, \qquad T_\chi(re^{it}) = re^{i(t + \chi(r))}$$

is a disk homeomorphism with inverse $T_{-\chi}$ and fixes the circle pointwise. Continuity at zero again follows from radius preservation. Conjugating by a closed-disk parametrization of V gives a nonidentity homeomorphism of $\overline{V}$ fixing ∂V. Extend it by the identity outside V; finite closed-set pasting proves continuity on the sphere. Composing the recognized model with this map changes its interior filling while preserving all prescribed exterior and boundary data. If there are no gaps, Corollary 14.4 explains why this particular freedom disappears.

Chapter 15

Limits and the general model problem

15.1 A limit of relations need not be a relation

The closedness of one fixed equivalence relation was essential throughout the construction. It controlled limits of its classes without changing the relation. A sequence of closed equivalence relations presents a different problem: its limit as a compact subset of $T \times T$ can fail transitivity. This can happen even when every approximating relation is unlinked and has only two nontrivial classes.

Example 15.1 (A lost transitive identification)**.** Put

$$a = 1, \qquad b = -1, \qquad c = -i, \qquad b_n = e^{i(\pi - 1/n)}, \qquad d_n = e^{i(\pi + 1/n)} \quad (n \geq 2).$$

Let R_n have the two non-singleton classes $\{a, b_n\}$ and $\{d_n, c\}$; all remaining points are singleton classes. Then R_n is closed and unlinked, but its Hausdorff limit as a subset of $T \times T$ is not transitive.

Proof. The two displayed classes are disjoint, so they define an actual equivalence relation. Its graph is the diagonal together with four off-diagonal pairs and is closed. The circular order of the four exceptional points is

$$a,\ b_n,\ d_n,\ c.$$

The pairs do not alternate, so they are unlinked. Singleton classes cannot alternate with any class. By Proposition 12.5 their filled hulls are disjoint.

The four exceptional ordered pairs converge respectively to $(a, b), (b, a), (b, c), (c, b)$. The diagonal remains fixed. Therefore the compact sets R_n converge in Hausdorff distance to

$$R_* = \Delta_T \cup \{(a, b), (b, a), (b, c), (c, b)\}.$$

To see that this is the entire limit, every point of R_n is either on the fixed diagonal or in that four-element list, and each listed limit has an approximating point. In R_* one has aR_*b and bR_*c, while $(a, c) \notin R_*$. Thus R_* is closed, reflexive and symmetric but not transitive.

This convergence is quantitative in the product metric $d((x, y), (x', y')) = \max\{|x - x'|, |y - y'|\}$. The exceptional moving coordinates satisfy

$$|b_n - b| = |d_n - b| = 2\sin(1/(2n)) \leq 1/n, \qquad d_H(R_n, R_*) \leq 1/n.$$

The diagonal points match themselves, and every one of the four remaining points has its explicitly matched limit. Both directed Hausdorff bounds therefore hold. This shows that the

transitivity defect persists even with a specified convergence rate; it is not caused by poor numerical resolution of the endpoints.

□

The corresponding limiting geodesics meet at b. They cannot remain two different classes of a decomposition, because a point cannot belong to two classes. Merging them forces a, b, c into one circle class. Our filled-hull convention then fills the entire hyperbolic triangle $H(\{a, b, c\})$, including its interior. This is an actual additional identification, not simply a different drawing of the same limiting two chords.

The example explains why convergence of a dense set of finite boundary data does not by itself produce a pinched-disk theorem. One must specify which new identifications are forced, verify that their hulls do not cross, and establish closedness after all required identifications. None of those tasks is replaced by a picture in which nearby endpoints appear to coincide.

Proposition 15.2 (A minimal closed equivalence completion)**.** *For any prescribed set of pairs $P \subset T \times T$, there is a unique smallest closed equivalence relation R_P containing P. If $f : T \to Y$ is continuous and Y is Hausdorff, then*

$$f(\zeta) = f(\xi) \text{ for all } (\zeta, \xi) \in P \quad \Longleftrightarrow \quad f \text{ is constant on every } R_P\text{-class.}$$

This construction alone does not ensure that R_P is unlinked.

Proof. Let $\mathcal{E}$ be the family of all closed equivalence relations on T containing P. It is nonempty because $T \times T$ is one such relation. Set

$$R_P = \bigcap_{E \in \mathcal{E}} E.$$

An intersection of closed sets is closed. Every E contains the diagonal, is symmetric, and is transitive. Hence their intersection has these properties: for the last one, if xR_Py and yR_Pz, the two pairs belong to every E, so (x, z) belongs to every E. Thus R_P is an equivalence relation containing P, and its defining intersection makes it the smallest one.

If f identifies every pair in P, its fiber relation is a closed equivalence relation by Proposition 11.1 or directly by pulling back the diagonal of $Y \times Y$. It belongs to $\mathcal{E}$, so it contains R_P. This proves one implication; the other follows from $P \subset R_P$.

Equivalently, the minimal completion has a universal property among continuous maps to Hausdorff spaces. If f identifies P, then

$$R_P \subset R_f, \qquad R_f = \{(x, y) : f(x) = f(y)\}, \qquad f = \bar{f} \circ q_P.$$

The factor $\bar{f} : T/R_P \to Y$ is continuous by the quotient property. Conversely any such factor identifies all pairs in P. This universal property concerns identifications in a topological quotient. It contains no assertion that the resulting classes have disjoint hyperbolic hulls; the crossing example below shows exactly why that is a further requirement.

For the final assertion, prescribe the two pairs $\{1, -1\}$ and $\{i, -i\}$. The relation with these two two-point classes and all other points singleton is already closed and is exactly the minimal completion. Its two classes alternate around the circle, so their diameters intersect. Closed equivalence completion has not repaired the crossing. □

One could impose unlinkedness as a separate requirement on a completion, but its existence, minimality and relation to the original geometric data would then require their own arguments. In particular, the preceding proposition is not a proof that a partially described lamination has a canonical cleaned version.

15.2 A precise conditional projection theorem

For a general full continuum K, source Theorem 9.29 describes a locally connected model using cutpoint landing data, limits of geodesics and a cleaning operation. Its density condition does not by itself state the final equivalence relation, the exact collapsed sets, or the continuity mechanism of the projection. The following conditional theorem separates the topological consequence of suitable data from the work required to produce those data.

Theorem 15.3 (A well-posed monotone model)**.** *Let $K \subset \mathbb{C}$ be a nonempty compact connected full set. Suppose $E \subset K \times K$ is a closed equivalence relation such that every class is connected and nonseparating in the sphere. Suppose in addition that the compact quotient K/E is locally connected. Extend E to the sphere by making every point outside K a singleton class. There are a sphere projection $p : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ and a locally connected full continuum $L \subset \mathbb{C}$ such that*

$$p^{-1}(L) = K, \qquad p(K) = L, \qquad p|_{\widehat{\mathbb{C}} \setminus K} : \widehat{\mathbb{C}} \setminus K \longrightarrow \widehat{\mathbb{C}} \setminus L \text{ is a homeomorphism.}$$

The fibers of $p|_K$ are exactly the E-classes, and $L \cong K/E$. The restriction $p|_K$ is a homeomorphism if and only if every E-class is a singleton.

Proof. The extended relation has graph

$$\widehat{E} = \Delta_{\widehat{\mathbb{C}}} \cup E.$$

The set E is compact, since it is closed in the compact product $K \times K$; consequently it is closed in the sphere product. Thus $\widehat{E}$ is closed. Its classes are the specified connected nonseparating compact sets and the remaining singletons. It is not the universal sphere relation, because K is bounded and points sufficiently near infinity are distinct singleton classes.

The inclusion of E in the sphere product causes no new limit pairs. If $(x_n, y_n) \in E$ converges in that product, compactness of K puts x, y in K, and closedness in K^2 then gives $(x, y) \in E$. The union formula also handles limits of singleton pairs:

$$(x_n, x_n) \to (x, y) \quad \Longrightarrow \quad x = y.$$

Mixed sequences can be checked by taking an infinite subsequence in one of the two closed pieces. Thus the ambient extension is closed for a proved reason, even when K itself is not locally connected.

Moore's theorem identifies the quotient sphere with a sphere. Choose the identification to send the class of ∞ to ∞, and let p be the resulting projection. Since K is saturated, $p^{-1}(p(K)) = K$. The restriction to the open saturated set $\widehat{\mathbb{C}} \setminus K$ is a homeomorphism onto its image: its fibers are singletons, and an open subset of this saturated singleton region is open and saturated in the whole sphere, so its quotient image is open.

Set $L = p(K)$. It is compact and connected and avoids infinity. Its complement is the connected image of $\widehat{\mathbb{C}} \setminus K$, so L is full. The map $p|_K$ has exactly the E-classes as fibers, hence compact quotient factorization identifies L with K/E. The assumed local connectivity of K/E is therefore exactly the local connectivity of L.

More precisely, if $q_K : K \to K/E$ is the quotient map, there is a homeomorphism $\theta : K/E \to L$ satisfying

$$p|_K = \theta \circ q_K.$$

The map θ exists because both sides identify exactly the same pairs; it is a continuous bijection between compact Hausdorff spaces. Local connectivity of L is transported through this homeomorphism from the stated hypothesis on K/E. There is no application of the compact-image theorem to K here, because the general source K was not assumed locally connected.

If every class is a singleton, $p|_K$ is a continuous bijection between compact Hausdorff spaces and is a homeomorphism. Conversely a homeomorphism is injective and cannot have a non-singleton fiber. □

There are three distinct hypotheses here. Closedness controls limits of fibers, nonseparation permits the ambient sphere quotient, and local connectivity of K/E controls the regularity of the resulting model. None is silently inferred from the others. A locally connected K can have many nontrivial monotone quotients; its local connectivity alone does not force an arbitrarily selected projection to be a homeomorphism. For a canonical construction with the stronger property

$$p|_K \text{ is a homeomorphism} \quad \Longleftrightarrow \quad K \text{ is locally connected},$$

one must also prove that its chosen relation is trivial whenever K is locally connected, and that it collapses something whenever K is not. The forward implication follows from local connectivity of the target, but the reverse implication is a statement about the construction, not about Moore's theorem alone.

The present theorem therefore records a fully proved route once the relation is available. It does not claim to close the general source assertion by replacing its unspecified construction with arbitrary equivalence classes. The source's particular density hypothesis and its proposed cleaning rule remain an identification and existence obligation beyond this conditional result.

15.3 A Cantor fan and its connected fibers

The radial example in source Exercise 9.30 is especially useful because a precise monotone-model question can be answered completely without guessing what an unspecified cleanup does.

Definition 15.4 (Cantor fan)**.** Let $X \subset T$ be a nonempty perfect totally disconnected compact set. Define

$$K_X = \{r\zeta : 0 \le r \le 1,\ \zeta \in X\}.$$

Here perfect means that no point is isolated, and totally disconnected means that every connected subset is a singleton. For a closed subset of the circle these hypotheses imply that X has empty interior.

Proposition 15.5 (Topology of the fan)**.** *The set K_X is a full continuum with empty interior. It is locally connected at 0 and not locally connected at any other point. Every connected subset of $K_X \setminus \{0\}$ lies on a single radial spoke. Its only cutpoint is 0.*

Proof. The map $(r, \zeta) \mapsto r\zeta$ is continuous on the compact space $[0, 1] \times X$, so its image is compact. Each segment $[0, \zeta]$ is connected and all these segments contain 0; their union is connected.

If $z \notin K_X$ and $|z| \le 1$, then $z \neq 0$ and its direction $z/|z|$ does not belong to X. The outward radial path from z to $2z/|z|$ therefore avoids K_X. Points with modulus greater than one already lie in the connected exterior of the closed unit disk. Joining such paths to the circle of radius two proves that the complement of K_X is path connected. Thus K_X is full.

Every disk centered at a nonzero point of K_X contains points with nearby directions outside X, because X has empty interior in T. Every disk centered at zero also contains a point of such a direction. Hence K_X has empty interior and equals its own boundary.

For $0 < \varepsilon < 1$, the relatively open set $K_X \cap B(0, \varepsilon)$ is a union of initial radial segments with common endpoint zero. It is connected. These sets give a connected neighborhood basis at zero.

On $K_X \setminus \{0\}$ the angular map

$$\alpha(z) = z/|z| \in X$$

is continuous. The image of a connected subset is connected and therefore a singleton. This proves the single-spoke assertion. If $x = r\zeta \neq 0$, choose a small neighborhood avoiding zero. Any connected subset of that neighborhood containing x lies on the spoke of ζ. But no such set is a neighborhood of x in K_X: by perfectness there are $\zeta_n \in X \setminus \{\zeta\}$ tending to ζ, and $r\zeta_n \to x$. Thus no connected open neighborhood of x can lie in the chosen neighborhood avoiding zero. Local connectivity fails at every nonzero point.

For an explicit angular continuity bound, if $|z - x| < r/2$ then $|z| > r/2$, and

$$\left|\frac{z}{|z|} - \frac{x}{|x|}\right| \leq \frac{|z - x| + \big||z| - |x|\big|}{|z|} \leq \frac{4|z - x|}{r}.$$

Thus the angular map is continuous on every region bounded away from the vertex. If a connected set in that region met two directions, its angular image would be a connected subset of X with at least two points. The single-spoke conclusion concerns arbitrary connected subsets, not only paths.

To state the neighborhood obstruction with quantifiers, fix $0 < \varepsilon < r/2$ and suppose that N were a connected neighborhood of x inside $K_X \cap B(x, \varepsilon)$. It avoids zero, so $N \subset \{t\zeta : 0 < t \leq 1\}$. Because it is a neighborhood there is $\delta > 0$ with $K_X \cap B(x, \delta) \subset N$. Perfectness gives $\xi \in X \setminus \{\zeta\}$ with

$$|\xi - \zeta| < \delta/r, \qquad |r\xi - x| = r|\xi - \zeta| < \delta.$$

This puts a point of another spoke in N, a contradiction. The proof excludes even connected neighborhoods, whether or not those neighborhoods are themselves open.

Removing zero disconnects K_X. For example, a separation of X into two nonempty relatively clopen sets pulls back under α to a separation of $K_X \setminus \{0\}$. For an explicit separation, choose distinct $\zeta, \xi \in X$. Each of the two circle arcs between them contains a point outside X, since X has empty interior. Choose one such point in each arc. The two components of the circle minus these chosen points meet X in nonempty complementary relatively open and closed sets, as required.

Now remove $x = r\zeta \neq 0$. The union A of all spokes with direction different from ζ, together with $\{t\zeta : 0 \leq t < r\}$, is connected, since all its pieces contain zero. Its closure relative to $K_X \setminus \{x\}$ is that entire space. Indeed, the only missing points lie farther out on the spoke of ζ, and they are limits of points with directions $\zeta_n \to \zeta$ distinct from ζ. The closure of a connected set is connected, also in a subspace. Therefore $K_X \setminus \{x\}$ is connected, and x is not a cutpoint.

For the density assertion, each missing point has the form $t\zeta$ with $t > r$. Choose distinct directions $\zeta_n \to \zeta$ in X. Then

$$t\zeta_n \in A, \qquad |t\zeta_n - t\zeta| = t|\zeta_n - \zeta| \to 0.$$

If $r = 1$, there are no missing points beyond the removed endpoint, so the conclusion is immediate. For $0 < r < 1$, the displayed sequence establishes closure inside the punctured space. Taking the closure in K_X instead would restore the deleted point x and would not be the connectedness argument that is required.

□

The connectedness of the punctured fan at a nonvertex point does not assert path connectedness. The density of the other spokes supplies a connected closure even though an individual outer piece of the punctured spoke has no radial path to the vertex. This distinction matters when translating a cutpoint test into a statement about legal arcs.

Theorem 15.6 (All monotone locally connected images collapse)**.** *Let Y be a compact Hausdorff locally connected space and $q : K_X \to Y$ a continuous surjection with connected fibers. Then Y is a singleton. Thus the point is the only locally connected monotone quotient of the Cantor fan.*

Proof. The image Y is connected because K_X is connected. Put $y_0 = q(0)$. Suppose there is a point $y \neq y_0$. Local connectivity supplies a connected open neighborhood V of y contained in $Y \setminus \{y_0\}$. Its preimage is a nonempty relatively open subset of K_X avoiding zero. It is also connected by Lemma 11.6: the map q is closed, because its domain is compact and its codomain Hausdorff, and its fibers are connected.

The role of each mapping hypothesis is visible if we name the preimage $W = q^{-1}(V)$:

$$W \neq \varnothing, \qquad W \text{ is open in } K_X, \qquad 0 \notin W, \qquad W \text{ is connected.}$$

Surjectivity gives the first assertion, continuity the second, exclusion of $q(0)$ the third, and the closed monotone inverse-image lemma the fourth. Connectedness is not inferred just from continuity, which ordinarily preserves images rather than inverse images. These four properties together contradict the fan's local single-spoke obstruction.

By Proposition 15.5, this connected preimage lies on one spoke. Yet a nonempty relatively open subset of K_X avoiding zero cannot lie on one spoke. If it contains $r\zeta$ with $r > 0$, openness and perfectness of X put $r\zeta_n$ in it for all sufficiently large n, for directions $\zeta_n \neq \zeta$ tending to ζ. This contradiction proves $Y = \{y_0\}$. □

In the category specified by the theorem, a finest locally connected monotone model is therefore a point: there is no nontrivial object to retain. The category is essential. The continuous radius map

$$\rho : K_X \to [0, 1], \qquad \rho(z) = |z|,$$

has a nontrivial locally connected image, but its fiber $\rho^{-1}(r) = rX$ is disconnected for $r > 0$. It is not a monotone map and does not contradict the theorem.

This resolves a precise interpretation of the Cantor-fan model question. Identifying it with a particular cleaned lamination in source Exercise 9.30 would additionally require proving that the source construction yields a monotone locally connected quotient with the stated meaning. We do not infer that missing identification merely from the word *model*. There is also no claim that an arbitrary fan automatically satisfies the separate angular-density hypothesis of source Theorem 9.29.

15.4 Locally connected finite fans can converge to this fan

Proposition 15.7 (A non-locally connected Hausdorff limit)**.** *There exist finite sets $X_n \subset X$ for which the finite fans K_{X_n} are locally connected full continua and*

$$d_{\mathrm{H}}(K_{X_n}, K_X) \longrightarrow 0.$$

More precisely, $d_{\mathrm{H}}(K_{X_n}, K_X) \leq d_{\mathrm{H}}(X_n, X)$.

Proof. Compactness gives a finite $1/n$-net $X_n \subset X$. For $r\zeta \in K_X$, choose $\xi \in X_n$ with $|\zeta - \xi| \leq 1/n$. Then

$$|r\zeta - r\xi| = r|\zeta - \xi| \leq 1/n.$$

The reverse Hausdorff inclusion is exact because $K_{X_n} \subset K_X$. Taking the infimum over admissible net radii gives the displayed bound.

A finite fan is compact and connected by the same union argument as before, and full by the same outward radial paths. At a nonvertex point, sufficiently small balls meet only its own spoke, so their relative intersections are interval neighborhoods. At zero, small balls cut connected initial segments from all its finitely many spokes. Thus the finite fan is locally connected everywhere.

The geometric estimate uses the same radial coordinate on both sets, rather than matching arbitrary nearest points in the plane. For any $\delta > d_H(X_n, X)$, every direction of X has a representative in X_n within δ, so

$$\sup_{z \in K_X} \operatorname{dist}(z, K_{X_n}) \le \delta, \qquad \sup_{z \in K_{X_n}} \operatorname{dist}(z, K_X) = 0.$$

Letting δ decrease to the angular Hausdorff distance proves the bound. Local connectivity of each approximant was established separately by its finite-spoke geometry; it is not a consequence of this distance estimate.

□

There is no common local-connectivity modulus for all these approximants. A concrete obstruction uses two nearby distinct directions. In a finite fan every path between points on different spokes must pass through zero: if it avoided zero, its connected image would have constant angular coordinate. Consequently a path joining two unit tips on different spokes has diameter at least one, however close those tips are. As increasingly fine nets acquire arbitrarily close distinct directions, a fixed small-path estimate must fail uniformly. The Hausdorff metric controls the location of the sets, but does not preserve the small connected neighborhoods required for local connectivity.

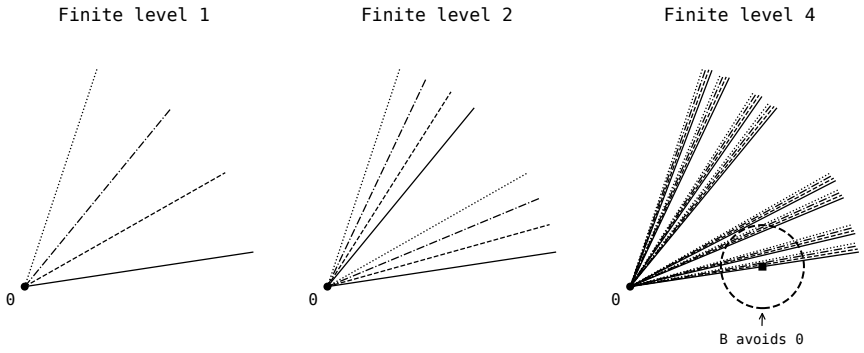


Figure 15.1: Finite stages of a Cantor-direction fan and a neighborhood away from its vertex. The small neighborhood meets distinct radial spokes which cannot be joined inside it. The drawings illustrate the single-spoke and perfectness arguments for the infinite fan; no finite drawing by itself proves failure of local connectivity.

15.5 Solved investigations on limits and quotients

Investigation 15.8 (Completing the missing triangle)**.** Determine the minimal closed equivalence completion of the limit relation R_* in Example 15.1.

Solution. Transitivity forces $a \sim c$ from $a \sim b$ and $b \sim c$. Thus every completion contains the class $\{a, b, c\}$. The relation with exactly this nontrivial class and all other points singleton is closed, contains R_*, and is an equivalence relation. Minimality therefore identifies it as the completion. Its hyperbolic hull is a filled triangle, since a, b, c are three distinct circle points and cannot be collinear. In contrast the set limit of the two original hulls is only $H(\{a, b\}) \cup H(\{b, c\})$. The additional filled region is forced by the completion convention.

Investigation 15.9 (Why nonseparation is a real hypothesis)**.** What fails if an equator of the sphere is made one class and all other points are singleton classes?

Solution. The relation is closed and all classes are connected, but the equator separates the sphere into two open hemispheres. In the quotient, removing the equator class leaves the disjoint images of those two hemispheres. The quotient is therefore not a sphere, since a sphere with one point removed is connected. The nonseparating-class hypothesis in Moore's theorem cannot be dropped even when only one class is nontrivial.

Investigation 15.10 (A locally connected source can collapse)**.** Give data satisfying Theorem 15.3 with locally connected K but noninjective projection.

Solution. Let $K = \overline{\mathbb{D}}$, and let $E = K \times K$ have one class. The class is connected and nonseparating in the sphere, and K/E is a point, hence locally connected. Extending the relation by singleton exterior points satisfies all hypotheses of the theorem. The resulting map collapses the entire disk. The source is locally connected but the projection on it is not a homeomorphism. A theorem asserting the converse must require a construction that keeps all classes trivial on locally connected input.

Investigation 15.11 (Radius is not monotone)**.** Compute the radius-map fibers and locate exactly why it escapes Theorem 15.6.

Solution. For $r = 0$ the fiber is $\{0\}$. For $0 < r \leq 1$ it is rX, homeomorphic to X under multiplication by r. The set X has more than one point and is totally disconnected, so these fibers are disconnected. The map is continuous and onto the locally connected interval, but the inverse-image connectedness lemma is unavailable. Indeed the inverse image of any nonempty open interval inside $(0, 1)$ consists of separated radial strips on the different spokes and is not connected.

Investigation 15.12 (Where perfectness is used)**.** Why does the point-quotient conclusion fail if X is a nonempty finite subset of the circle?

Solution. A finite fan is locally connected by Proposition 15.7; its identity map has connected singleton fibers and a nontrivial locally connected image. At a nonvertex point, a sufficiently small relative neighborhood lies entirely on its one spoke. Thus the contradiction in the proof of Theorem 15.6 no longer applies. Perfectness supplied other directions arbitrarily close to every direction. Finiteness removes exactly that property, while leaving compactness and fullness intact.

Investigation 15.13 (A quantitative failure of uniform local connectivity)**.** For $F_n = [0, 1] \cup [0, e^{i/n}]$, show that no common small-path modulus can make these finite fans uniformly locally connected.

Solution. Each F_n is a locally connected full continuum. Its two unit tips satisfy

$$|1 - e^{i/n}| = 2\sin(1/(2n)) \leq 1/n \longrightarrow 0.$$

Every connected path in F_n joining those tips passes through zero, because removing zero leaves two components. Its image therefore has diameter at least $|1 - 0| = 1$. For the target diameter $\varepsilon = 1/2$, every proposed common $\delta > 0$ fails once $1/n < \delta$: the endpoints are closer than δ but cannot be joined by a path of diameter less than ε. The limiting set is a single segment and is locally connected; this example shows that even a regular limit does not rescue a uniform modulus for the approximants.

The finite-fan limit and the failed relation limit answer different questions. The first concerns regularity of the underlying continuum, the second the algebra of identifications. A general

pinched-model construction must control both. The locally connected recognition theorem does so through an actual closed unlinked boundary relation and a null-family pasting argument. For the general source construction, those checks must be supplied for its specified limiting and cleaning procedure before its projection and rigidity claims can be recorded as proved.

Chapter 16

Quantitative refinement and geometric stopping rules

Source correspondence. Lyubich, Sections 9.1.3 and 9.4: shrinking puzzle pieces and finite descriptions of a boundary quotient. The estimates below require the additional metric hypotheses stated explicitly.

A finite description of a puzzle contains several kinds of information. An address records which pieces contain a point; an angular partition records which accesses are identified; a drawing supplies approximate positions. These records answer different questions. In particular, increasing the length of an address does not by itself decrease a Euclidean error. This chapter specifies the extra information that turns refinement into a stopping rule. Its elementary estimates apply equally to compact subsets of the plane and to compact metric spaces, and their limitations are useful when the topology of the underlying puzzle remains difficult.

16.1 A refinement tree with a certified mesh

Definition 16.1 (Compact refinement data)**.** Let K be a nonempty compact metric space. A refinement tree consists of finite nonempty levels T_n, with $T_0 = \{o\}$, and a parent map from T_{n+1} onto T_n. To each vertex v attach a nonempty compact set $K_v \subset K$, with $K_o = K$, such that

$$K_w \subset K_v \quad (w \text{ a child of } v), \qquad K_v = \bigcup_{w \text{ a child of } v} K_w.$$

Different children may overlap. The level mesh is

$$\delta_n = \max_{v \in T_n} \operatorname{diam} K_v.$$

An address is an infinite compatible sequence $(v_n)_{n \geq 0}$.

The requirement that the levels are finite makes the displayed maximum literal. It also means that a complete level is a finite certificate only when the containment, covering and diameter assertions themselves have been established. A list of sampled vertices is not such a certificate. No disjointness assumption is needed for the first result.

Proposition 16.2 (Localization by mesh)**.** *Suppose $\delta_n \to 0$. Every address has a unique represented point $\pi(v_0, v_1, \ldots) \in K$, every point of K has at least one address, and any choice $x_n \in K_{v_n}$ satisfies*

$$d(x_n, \pi(v_0, v_1, \ldots)) \leq \delta_n.$$

Two addresses agreeing through level n represent points at distance at most δ_n. Their later symbols need not agree.

Proof. The sets along an address are nested, nonempty and compact. Any finite intersection is the last of its members, so compactness gives a point in their full intersection. If x, y lie in that intersection, then

$$0 \le d(x,y) \le \operatorname{diam} K_{v_n} \le \delta_n \quad \text{for every } n;$$

letting n increase proves $x = y$. The same inequality with x_n proves the approximation estimate. Given $x \in K$, start at the root and choose a child containing x at each step, which is possible by the covering identity. This produces an address of x. Finally, points represented by two addresses with the same nth vertex both lie in that compact set, and its diameter gives the asserted bound.

One can make the compactness of the address description explicit as well. Give each finite level T_n the discrete topology. A sequence of addresses has a subsequence with constant first coordinate, a further subsequence with constant second coordinate, and so on. On the diagonal subsequence, each fixed coordinate is eventually constant. Its limiting coordinates remain compatible because each parent equation already holds eventually:

$$\operatorname{parent}(v_{n+1}^{(j)}) = v_n^{(j)} \quad \Longrightarrow \quad \operatorname{parent}(v_{n+1}) = v_n.$$

In the address metric introduced below this is convergence: agreement through level n makes the distance at most ρ^n. Thus the space of compatible addresses is sequentially compact and, being metric, compact. The represented-point map is continuous by the mesh estimate, since

$$a^{(j)} \to a \quad \Longrightarrow \quad \forall n\ \exists j_n\ \forall j \ge j_n : d\big(\pi(a^{(j)}), \pi(a)\big) \le \delta_n.$$

First choose n from a desired error tolerance, then choose j_n from coordinate agreement. This order of choices proves continuity without assuming an injective coding or an effective formula for the modulus. For surjectivity, the successive choices of a child containing a given x can be made in a fixed order on each finite level. This removes any ambiguity in the existence construction while still allowing different orders to produce different addresses of a boundary point.

□

Here the metric information is an upper bound on an entire set. A curve connecting its sampled corners may leave the sampled region, and a thin finger may have a distant endpoint. Consequently the maximum distance between displayed corner samples cannot replace $\operatorname{diam} K_v$ unless an additional enclosure theorem supplies that replacement.

Proposition 16.3 (Representatives and approximate representatives)**.** *Choose $x_v \in K_v$ for every $v \in T_n$ and set $S_n = \{x_v : v \in T_n\}$. Then the Hausdorff distance satisfies $d_H(K, S_n) \le \delta_n$. If approximate representatives $\widetilde{x}_v$ in an ambient metric space obey $d(x_v, \widetilde{x}_v) \le \eta_n$, then*

$$d_H(K, \widetilde{S}_n) \le \delta_n + \eta_n, \qquad \widetilde{S}_n = \{\widetilde{x}_v : v \in T_n\}.$$

Proof. For every $x \in K$, a level-n piece containing x supplies a representative within δ_n. Conversely, every exact representative belongs to K. With approximate representatives the first directed distance is at most $\delta_n + \eta_n$ by the triangle inequality, and the second is at most η_n. Taking the larger directed distance proves the formula.

Here are the two directed estimates before taking that maximum:

$$\sup_{x\in K} \inf_{v\in T_n} d(x, \widetilde{x}_v) \le \sup_{x\in K} \inf_{v : x\in K_v} \big(d(x, x_v) + d(x_v, \widetilde{x}_v)\big) \le \delta_n + \eta_n,$$
$$\sup_{v\in T_n} \inf_{x\in K} d(\widetilde{x}_v, x) \le \sup_{v\in T_n} d(\widetilde{x}_v, x_v) \le \eta_n.$$

The restricted infimum in the first line is over a nonempty finite set, by coverage. In the second line the particular admissible point x_v is used; there is no need to compute a nearest point of

K. Repeated representatives do not affect either infimum, so distinct vertices need not produce distinct stored points. Conversely, if a numerical routine omits a vertex without establishing that its piece is covered by the remaining pieces, the first restricted infimum may become empty. The proof shows exactly where that omission would invalidate the enclosure.

□

This conclusion is a geometric enclosure of a compact set. It says nothing about the number of complementary components, local connectivity, or the continuity of a function defined on the set. Hausdorff approximation and topological identification require separate arguments.

16.2 Explicit moduli and address metrics

Suppose reference pieces Q_v lie in a compact metric space X and $K_v = F(Q_v)$ for a map $F : X \to Y$. A supplied modulus of continuity is a nondecreasing function $\omega : [0, \infty) \to [0, \infty)$ with $\omega(t) \to 0$ as $t \downarrow 0$ and

$$d_Y(F(x), F(y)) \le \omega(d_X(x, y)) \qquad (x, y \in X).$$

If $\operatorname{diam} Q_v \le h_n$, direct substitution gives

$$\operatorname{diam} K_v \le \omega(h_n), \qquad \delta_n \le \omega(h_n). \tag{16.1}$$

Uniform continuity ensures that some modulus exists on a compact domain. It does not reveal numerical constants or a computable rate from a picture. The stopping criterion is therefore conditional on having the displayed modulus as justified input.

Proposition 16.4 (Two explicit stopping bounds)**.** *Let $h_n \le Db^{-n}$, where $D > 0$ and $b > 1$. If $\omega(t) = Lt^{\alpha}$ for $L > 0$ and $0 < \alpha \le 1$, it is enough to take*

$$n \ge \max\left\{0, \left\lceil \frac{\log(LD^{\alpha}/\varepsilon)}{\alpha \log b} \right\rceil\right\}$$

to obtain $\delta_n \le \varepsilon$. If instead the certified bound on these scales is

$$\omega(t) \le \frac{A}{(\log(eD/t))^{\beta}} \quad (0 < t \le D), \qquad A, \beta > 0,$$

it is enough to take

$$n \ge \max\left\{0, \left\lceil \frac{(A/\varepsilon)^{1/\beta} - 1}{\log b} \right\rceil\right\}.$$

Proof. In the first case (16.1) gives $\delta_n \le LD^{\alpha}b^{-\alpha n}$. Taking logarithms of this positive expression gives exactly the first inequality. In the second case, monotonicity of the proposed upper bound and $\log(eD/(Db^{-n})) = 1 + n \log b$ give

$$\delta_n \le A(1 + n \log b)^{-\beta}.$$

Solving this inequality for n gives the second formula. The outer maximum handles tolerances already satisfied at level zero.

To check the direction of every inequality, put $H = LD^{\alpha}$ and assume first $0 < \varepsilon < H$. Since $\alpha \log b > 0$,

$$Hb^{-\alpha n} \le \varepsilon \iff -\alpha n \log b \le \log(\varepsilon/H) \iff n \ge \frac{\log(H/\varepsilon)}{\alpha \log b}.$$

For a real number s, the least integer satisfying $n \geq s$ is $\lceil s \rceil$. When s is already an integer there is no extra level: equality in the certified error bound is allowed. If $\varepsilon \geq H$, level zero already passes, which explains the outer maximum rather than an unexplained negative depth. For the logarithmic modulus and $0 < \varepsilon < A$, the corresponding positive-power calculation is

$$\frac{A}{(1 + n \log b)^{\beta}} \leq \varepsilon \iff 1 + n \log b \geq (A/\varepsilon)^{1/\beta}.$$

Taking a positive βth root preserves order. These are least depths for the displayed upper-bound sequences, not necessarily least depths for the actual pieces, whose diameters may be much smaller. That distinction also explains why a sufficient stopping calculation is not an assertion of optimal computational cost.

□

There is no assertion here that the boundary extension of an arbitrary Riemann map has a Hölder modulus. A weak logarithmic modulus, when it is the available input, can demand many more levels than a geometric rate. The same formal address tree therefore supports very different accuracy costs under different geometric realizations.

16.2.1 From a conformal collar to an explicit mesh

The derivative estimate in Proposition 4.3 provides one concrete source of the mesh input. Fix a nonsingular exterior annulus $r \leq |w| \leq R$ with $1 < r < R$, and suppose a certified constant M bounds $|\Phi'_K|$ throughout that annulus. Divide the radial interval into 2^n equal parts and the angle circle into 2^n equal parts. For $n \geq 1$, each closed polar cell has radial width $(R - r)2^{-n}$ and angular width 2^{-n} turns. Hence its image has diameter at most

$$\delta_n \leq M(R - r + 2\pi R)2^{-n}. \tag{16.2}$$

These cells cover the compact annulus, and each coarse cell is the union of four cells at the next level. Their images therefore form genuine compact refinement data, with four children per cell after the first level. A single image point may still belong to two adjacent closed cells.

Let $C = M(R - r + 2\pi R)$. A tolerance $\varepsilon < C$ is achieved at $n = \lceil \log_2(C/\varepsilon) \rceil$, bounded below by one. There are 4^n polar cells, so complete enumeration has an upper bound

$$4^n < 4(C/\varepsilon)^2$$

when that ceiling is at least one. This estimate counts cells rather than distinct image samples, and the factor C explicitly depends on the chosen annulus and conformal derivative bound. If the tolerance already permits level one, four cells suffice regardless of the displayed asymptotic expression. The distinction between a single addressed point and an entire two-parameter collar explains the different storage demands.

To let the inner radius approach one, write $r = r_n$ and, if justified, replace M by M_n. The same argument then gives a bound proportional to $M_n 2^{-n}$, not automatically to 2^{-n}. For example, a known estimate $M_n \leq A2^{\theta n}$ with $\theta < 1$ gives geometric decay $O(2^{-(1-\theta)n})$ when the outer radii stay bounded. At $\theta = 1$ this particular estimate need not decay, and for $\theta > 1$ it grows. Failure of the estimate to decay does not prove that the image cells fail to shrink; it proves that this chosen derivative bound has ceased to certify their shrinking. A boundary modulus, when available, can provide a different route through Proposition 4.2. This is a useful example of separating failure of a certificate from failure of the underlying geometric statement.

16.2.2 Nonuniform contraction along an address

A uniform geometric factor is sufficient but not necessary. Suppose every parent-child pair satisfies a proved bound

$$\operatorname{diam} K_w \le q_w \operatorname{diam} K_v, \qquad 0 < q_w \le 1.$$

Along an address (v_n), induction gives

$$\operatorname{diam} K_{v_n} \le \operatorname{diam} K \prod_{j=1}^{n} q_{v_j} = \operatorname{diam} K \exp\left(-\sum_{j=1}^{n} \log(1/q_{v_j}) \right). \tag{16.3}$$

Thus divergence of the nonnegative sum in the exponent is a sufficient condition for a singleton impression on that address. A sufficient stopping test is that the partial sum reaches $\log(\operatorname{diam} K/\varepsilon)$ when $\operatorname{diam} K > \varepsilon > 0$. This test allows occasional noncontracting steps with $q_w = 1$. It requires the inequality for every relevant child, not an average contraction estimated from selected successful paths.

The logarithmic form also gives a useful comparison with linear losses. For $0 < q \le 1$, integration of $1/t$ over $[q, 1]$ gives

$$1 - q \le -\log q \le \frac{1-q}{q}.$$

Consequently $\sum_j (1 - q_{v_j}) = \infty$ suffices to force the certified product to zero. If the factors are all at least a fixed $q_* > 0$, the two sums diverge together, because the upper bound is then at most $(1 - q_{v_j})/q_*$. No lower bound on the actual diameters follows from these comparisons; they concern the multiplicative upper certificate. An exact nongeometric example is

$$q_{v_j} = \frac{j}{j+1}, \qquad \prod_{j=1}^{n} q_{v_j} = \frac{1}{n+1}, \qquad \sum_{j=1}^{n} -\log q_{v_j} = \log(n+1).$$

Each factor approaches one, yet the telescoping product tends to zero. If instead a factor $1/2$ occurs only at indices $2, 4, 8, \ldots$ and all other factors equal one, the product at depth $n \ge 2$ is $2^{-\lfloor \log_2 n \rfloor}$. Long stretches without contraction are compatible with eventual localization. The explicitly counted contracting indices, rather than the address length alone, determine the stopping bound.

There is also a useful uniformity principle. Fix $\varepsilon > 0$ and suppose every infinite address eventually reaches a vertex whose certified diameter bound is at most ε. If the tree is finitely branching and these bounds never increase along a branch, then there is one finite depth by which all addresses have reached such a vertex. To prove this, suppose otherwise. The vertices whose ancestor bounds all exceed ε form a subtree of arbitrarily large depth. Starting at its root, choose a child with descendants at arbitrarily large depths; one exists because there are only finitely many children. Repeat this choice to construct an infinite address consisting entirely of failing vertices, contradicting the hypothesis. This is a direct finite-branching argument, with no appeal to an observed termination rate.

The principle proves existence of a common finite depth for each tolerance but does not provide its value. Formula (16.3) provides a numerical depth only if the cumulative contraction sums are themselves bounded below by known quantities. Pointwise divergence along all addresses and an explicit usable uniform rate remain distinct inputs.

Fix $0 < \rho < 1$. For two distinct infinite addresses, let ℓ be the largest common level and set $d_\rho = \rho^\ell$; for identical addresses set $d_\rho = 0$. This is an ultrametric: two addresses each agreeing

with a third through level m agree with one another through level m. If $\delta_n \le Cq^n$ for $0 < q < 1$, Proposition 16.2 yields

$$d(\pi(a), \pi(b)) \le C d_\rho(a, b)^\gamma, \qquad \gamma = \frac{\log q}{\log \rho} > 0. \tag{16.4}$$

The inequality remains valid for identical addresses by taking the limit. It proves continuity of the coding map and quantifies it. It does not make that map injective: distinct addresses can encode a shared boundary point.

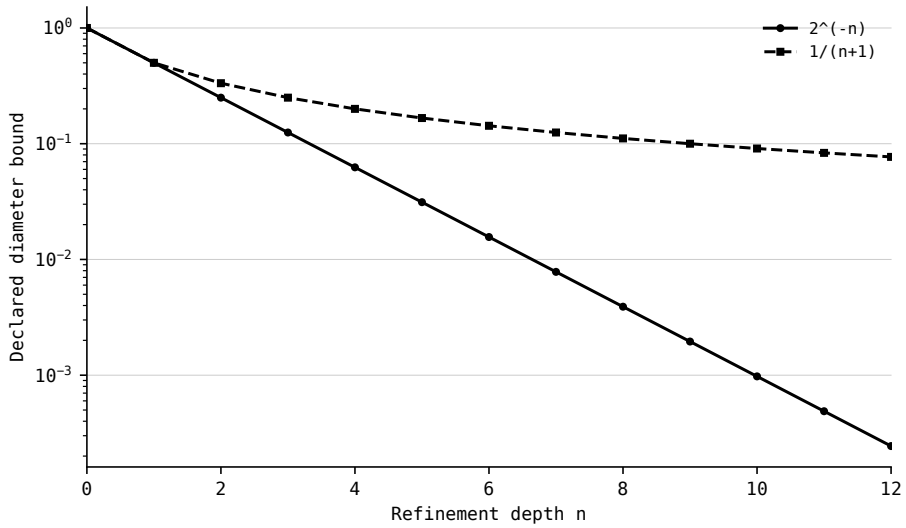


Figure 16.1: The certified bounds 2^{-n} and $(1+n)^{-1}$ for integer refinement levels $n \ge 0$. These are exact model bounds, not measurements of a particular conformal puzzle.

16.3 Adaptive refinement and what topology still needs

Proposition 16.5 (Finite termination with a uniform rate)**.** *Assume each vertex has finitely many children and* $\operatorname{diam} K_v \le Cq^{|v|}$*, with* $0 < q < 1$*. Refine every active piece whose certified diameter upper bound exceeds* ε*, using the bound* $Cq^{|v|}$*. The process stops after finitely many refinements. Its terminal pieces cover* K *and each has diameter at most* ε*.*

Proof. Choose N with $Cq^N \le \varepsilon$. No vertex at depth N needs to be refined, and a finitely branching tree has only finitely many vertices through a fixed depth. Replacing a piece by all its children preserves the union, so the final union is K. The stopping test gives the final diameter assertion. The proof concerns the stated bound: an implementation that keeps using a needlessly large bound may fail its stopping test even when the actual pieces are small.

More explicitly, the active vertices always form an antichain: no active vertex is an ancestor of another, because a vertex is removed at the moment its children are inserted. Every infinite address meets the active set exactly once, and the corresponding pieces cover K. Each refinement preserves both statements. If $C = 0$, the root has diameter zero and already passes. For $C > 0$, put

$$N = \max\left\{0, \left\lceil \frac{\log(C/\varepsilon)}{\log(1/q)} \right\rceil \right\},$$

every active vertex has depth at most N. If, in addition, each vertex has at most $b \ge 2$ children, the total number of vertices ever created is bounded by

$$1 + b + \cdots + b^N = \frac{b^{N+1} - 1}{b - 1}, \qquad \#\{\text{terminal active vertices}\} \le b^N.$$

For the second inequality, extend each active vertex to depth N; different antichain vertices have disjoint sets of such descendants, and each has at least one because the parent maps are onto. This counts vertices rather than distinct physical pieces and therefore tolerates overlap. Without a uniform branching bound, the finite-level argument still proves termination, but the displayed numerical work bound is no longer available. If there is only one child at every step, the separate bound is $N + 1$ created vertices and one terminal piece.

□

To obtain a connectivity conclusion at x from a shrinking family, one must know that the relevant connected subsets actually furnish relative neighborhoods of x. More precisely, if for each $\varepsilon > 0$ there is a connected relative neighborhood N of x in K, with $N \subset B(x, \varepsilon) \cap K$, then K is weakly locally connected at x, in the terminology of the topology preliminaries. If this holds at every point, Proposition 3.1 gives local connectivity of K. For local connectivity at one selected point, connected relatively open neighborhoods suffice directly. A family of small connected sets merely containing x is insufficient. The singleton $\{x\}$ would otherwise prove local connectivity of every space. In the normalized puzzle arguments of this volume, the explicitly admissible or perfect connected pieces in Definition 5.5, with the shrinking hypotheses of Proposition 7.2, supply the needed neighborhood information. The broad source cut-curve definition alone and a diameter routine do not supply it.

16.4 Solved investigations

Investigation 16.6 (A Hölder stopping calculation)**.** Reference pieces have diameter at most $4 \cdot 2^{-n}$ and their images satisfy $d(Fx, Fy) \le 3d(x, y)^{1/2}$. Find a sufficient depth for error $1/100$, and include representative error at most $1/1000$.

Solution. The mesh bound is $6 \cdot 2^{-n/2}$. For exact representatives it suffices that $2^{n/2} \ge 600$, hence $n \ge 19$, since $2^9 < 600 < 2^{19/2}$. With approximate representatives and a total tolerance $1/100$, reserve $9/1000$ for the mesh. The condition becomes $2^{n/2} \ge 2000/3$. Depth 19 still suffices because $2^{19} = 524288 > (2000/3)^2$. Proposition 16.3 then gives $d_H(K, \widetilde{S}_{19}) \le 6 \cdot 2^{-19/2} + 10^{-3} < 10^{-2}$. This is an exact inequality, not a decimal estimate of a measured error.

Investigation 16.7 (Two binary descriptions of one point)**.** Use the closed dyadic subintervals of $[0, 1]$. Exhibit two addresses of $1/2$ and explain why the coding estimate does not imply uniqueness of addresses.

Solution. At level one choose either $[0, 1/2]$ or $[1/2, 1]$. In the left interval always choose the right child; in the right interval always choose the left child. The resulting intervals have common intersection $\{1/2\}$. Their lengths are 2^{-n}, and their binary labels are $0.0111\ldots$ and $0.1000\ldots$. Formula (16.4) bounds distances between represented points from above. It gives no positive lower bound for distinct symbolic strings and therefore permits these equal represented points.

Investigation 16.8 (A refinement with no geometric progress)**.** Construct a binary refinement tree whose address length tends to infinity while its mesh remains one. Explain which hypothesis of Proposition 16.2 fails.

Solution. Attach $K_v = [0, 1]$ to every vertex of the full binary tree. Every child is contained in its parent and the children cover it, but $\delta_n = 1$ for all n. Every address has impression $[0, 1]$ rather than a singleton. Thus the containment and covering tests alone pass while the mesh hypothesis fails. The construction also shows why counting distinct labels, without checking their geometric meaning, supplies no accuracy certificate.

Investigation 16.9 (Arbitrarily slow single-address localization)**.** Let $a_n > 0$ decrease to zero. Give a nested sequence realizing exactly this diameter rate, and show why no predetermined geometric rate follows from singleton intersection alone.

Solution. Take $P_n = [0, a_n]$. Then $\operatorname{diam} P_n = a_n$ and $\bigcap_n P_n = \{0\}$. In particular $a_n = 1/(n+1)$ cannot satisfy $a_n \le Cq^n$ for fixed C and $q < 1$: the ratio $q^{-n}/(n+1)$ tends to infinity, as follows, for example, by comparing successive ratios with a number strictly larger than one for large n. This example is a single puzzle nest, not a complete covering tree. It addresses precisely the inference from a singleton impression to a uniform or geometric shrinking rate.

Investigation 16.10 (A certified outer enclosure)**.** Suppose finitely many compact pieces cover K, each has diameter at most δ, and each chosen representative is known within η. Show that the union of closed balls of radius $\delta + \eta$ about the approximate representatives contains K. Can its connected components be identified with the components of K?

Solution. For $x \in K$, choose a containing piece and its exact representative s. Then $d(x, \widetilde{s}) \le d(x, s) + d(s, \widetilde{s}) \le \delta + \eta$. This proves the enclosure. Component identification fails: take $K = \{0, 3\delta/2\} \subset \mathbb{R}$, exact representatives, and allow the diameter bound $\delta > 0$ for its singleton pieces. The two balls of radius δ overlap, so the enclosure is connected although K has two components. An upper diameter bound need not be sharp enough to preserve topology, even when it is completely correct.

Chapter 17

Exact circular order and finite noncrossing certificates

Source correspondence. Lyubich, Section 9.4: unlinked angle classes and their chordal realizations. Only finite, explicitly supplied partitions are considered in this chapter.

An angular drawing is particularly vulnerable to rounding near a shared endpoint or a narrow gap. Fortunately, the combinatorial question of whether two chords cross does not require evaluating a sine, a cosine or an approximate intersection point. It depends on circular order alone. This observation leads to exact certificates for finite data, while also making clear why finite data cannot certify an entire infinite lamination. Throughout the chapter an angle is measured in turns:

$$t \in \mathbb{R}/\mathbb{Z}, \qquad z(t) = \exp(2\pi i t).$$

The canonical representative lies in $[0, 1)$.

17.1 The four-endpoint test

For distinct angles a, b, write $(a, b)_{\mathrm{ccw}}$ for the open counterclockwise arc from a to b. Define the Boolean predicate

$$J(a, b; t) = \begin{cases} a < t < b, & a < b, \\ (t > a) \text{ or } (t < b), & a > b. \end{cases}$$

The comparisons use canonical representatives. Endpoints are excluded. The notation on the right means the truth value of the indicated statement.

Proposition 17.1 (Crossing criterion)**.** *Suppose a, b, c, d are four distinct points of $\mathbb{R}/\mathbb{Z}$. The open straight chords joining $z(a)$ to $z(b)$ and $z(c)$ to $z(d)$ intersect if and only if*

$$J(a, b; c) \neq J(a, b; d).$$

Equivalently, the two endpoint pairs alternate in circular order. If two nonidentical chords share one endpoint, their interiors do not intersect. Identical chords overlap and must be reported separately.

Proof. The straight line through $z(a)$ and $z(b)$ divides the remaining points of the circle into the two open arcs from a to b, lying in opposite open half-planes. This can be checked after a rotation taking the midpoint angle of one arc to zero: the line becomes $\operatorname{Re} z = \cos\theta$ for some $0 < \theta < \pi$, and the sign of $\cos t - \cos\theta$ is constant on each of the two arcs. Thus $z(c)$ and $z(d)$ lie in opposite half-planes exactly when the displayed Boolean values differ.

In that case their segment crosses the line at an interior point. Since the open segment joining distinct circle points lies inside the open disk, this crossing lies in the intersection of the line with the disk, which is exactly the open chord from $z(a)$ to $z(b)$. If both endpoints lie in the same open half-plane, the segment stays there and cannot cross the line. This proves the criterion. Chords with a common endpoint lie on distinct straight lines unless they coincide: a line meets the circle in at most two points. Distinct lines intersect only at that common endpoint, proving the remaining assertion.

The crossing step admits an explicit affine calculation. In the rotated coordinate write $\ell(z) = \Re z - \cos\theta$, and suppose $\ell(z(c)) > 0 > \ell(z(d))$. The unique zero along their segment occurs at

$$z_* = (1-s)z(c) + sz(d), \qquad s = \frac{\ell(z(c))}{\ell(z(c)) - \ell(z(d))} \in (0,1).$$

The zero is unique because the denominator is nonzero. Moreover, for any two distinct unit vectors u, v and $0 < s < 1$,

$$|(1-s)u + sv|^2 = 1 - s(1-s)|u-v|^2 < 1.$$

Thus z_* is strictly inside the disk, not at an endpoint of its intersection with the line. Those two endpoints are precisely $z(a)$ and $z(b)$, so z_* belongs to the open first chord as well. The same calculation shows why four distinct endpoints are essential: if an endpoint is shared, the affine zero can occur at $s = 0$ or $s = 1$, which is a different geometric event and is excluded from the open-chord test. This derivation establishes the predicate without numerically evaluating s; exact circular order already determines the signs it needs.

□

The proof handles diameters as well as short chords. It does not assign a crossing to a tangency produced by a rounded drawing: the exact question first distinguishes four distinct endpoints, one shared endpoint, and an identical unordered pair. In a partition different blocks cannot share an angle, but a list of geometric edges may legitimately have repeated vertices within a block.

17.2 Rational input and independently readable witnesses

Let an input angle be represented as a rational number p/q, with $q > 0$ and $0 \le p < q$. Reduction to lowest terms is convenient but not essential for comparisons, since

$$\frac{p}{q} < \frac{r}{s} \iff ps < rq, \qquad \frac{p}{q} = \frac{r}{s} \iff ps = rq. \tag{17.1}$$

Both tests use integers. A proof of the first equivalence simply multiplies by the positive number qs; this positivity is why zero or negative denominators are rejected or normalized before the test.

Definition 17.2 (Finite partition certificate)**.** The input comprises distinct rational angles $t_1, \dots, t_N$ and a list of nonempty blocks whose disjoint union is the index set $\{1, \dots, N\}$. It is noncrossing when no four distinct angles a, c in one block and b, d in a different block occur in the cyclic order a, b, c, d. This is a definition for an actual set partition; a merely symmetric pairwise compatibility relation is not accepted as a partition.

Proposition 17.3 (A finite exact decision procedure)**.** *For the data in Definition 17.2, the following procedure decides noncrossing using integer arithmetic. First check the partition axioms and distinctness of the angles. Then test every chord whose two endpoints belong to one block against every such chord from another block, using Proposition 17.1. Accept exactly when all tests report no crossing. A rejection for crossing has a witness consisting of two block indices and four endpoint indices.*

Proof. If a test finds a crossing, the four endpoints are in two different blocks and therefore distinct. Their circular order alternates by Proposition 17.1; this is a forbidden quadruple in the definition. Conversely, a forbidden quadruple supplies a pair of within-block chords which the exhaustive procedure tests, and the same proposition makes that test fail. Each truth value is determined by finitely many exact comparisons (17.1). Thus the procedure terminates and its acceptance is equivalent to the stated finite property. □

There are at most $N(N-1)/2$ within-block chords, so a deliberately simple implementation uses $O(N^4)$ chord-pair tests. This is an upper bound for the complete method, not a claim of optimality. If every numerator and denominator has at most B binary digits, grade-school multiplication and comparison take $O(B^2)$ bit operations per rational comparison. The direct bound is consequently

$$O(N^4 B^2) \tag{17.2}$$

bit operations for the crossing portion. Storage can remain $O(NB+N)$ bits apart from $O(N \log N)$ index storage if pairs are generated successively; storing all pairwise answers is unnecessary. The input validation and index storage must of course also be included in an implemented cost account.

The finite bound can be audited before implementing any geometry. If the block sizes are $n_1, \dots, n_k$, let

$$M = \sum_{i=1}^{k} \binom{n_i}{2} \le \binom{N}{2}, \qquad \#\{\text{tested unordered chord pairs}\} \le \binom{M}{2} \le \frac{N^2(N-1)^2}{8}.$$

Pairs from the same block and pairs sharing an endpoint are discarded according to the input contract, so this deliberately loose bound remains valid. A positive integer of at most B bits is smaller than 2^B; each cross-product in a rational comparison is smaller than 2^{2B}, and their signed difference needs at most $2B+1$ bits. Grade-school multiplication forms at most B^2 bit-products, followed by bounded carry operations of the same asymptotic cost. No denominator product is ever converted to a floating-point number. After sorting and recording ranks, one obtains the alternative bound

$$O(N \log N\, B^2 + N^4 \log N)$$

for sorting plus crossing comparisons in bit operations. The rank table requires $O(N \log N)$ bits. The original rational input still costs $O(NB)$ bits, so ranking changes the repeated-comparison cost rather than making that input disappear. A rejected crossing needs only the four endpoint indices and two block indices; a checker then repeats a constant number of exact comparisons instead of trusting the search order.

A useful alternative sorts the angles once by exact comparison and replaces each endpoint by its rank. All later crossing tests then concern integers between 1 and N. The initial sort can use $O(N \log N)$ rational comparisons; subsequent tests use order on $O(\log N)$-bit ranks. Neither approach evaluates the positions $z(t_j)$. The resulting certificate is invariant under every orientation-preserving homeomorphism of the circle, since such a map preserves cyclic order. Its verdict is about the supplied partition, not about the correctness of an unproved claim that these blocks are boundary fibers of a given hull.

17.3 A faster certificate for a perfect matching

Assume now that every block has size two, so each of the $2m$ endpoints belongs to exactly one chord. Cut the circular order at a point which is not an endpoint and list endpoints in increasing order. The first endpoint of a chord is its opening occurrence and the second its closing occurrence. A stack stores unmatched openings: append at an opening and remove the last stored entry at a closing.

Proposition 17.4 (Stack criterion for a matching)**.** *A perfect matching of distinct circular endpoints is noncrossing if and only if every closing occurrence matches the last unmatched opening. After sorting, this test uses* $2m$ *stack updates and* $O(m)$ *stored indices.*

Proof. Suppose a chord A closes while the most recent unmatched opening belongs to another chord B. The order of those four endpoints is

$$A_{\text{open}} < B_{\text{open}} < A_{\text{close}} < B_{\text{close}},$$

since B is still unmatched. The chords cross by Proposition 17.1. Hence noncrossing forces the stack rule. Conversely, if two chords cross, their endpoints alternate in the linear order after possibly exchanging their names, exactly as displayed. When A closes, the opening of B remains above that of A in the stack. The stack rule therefore fails at that time, or has already failed at an earlier closing. If the rule never fails, no crossing exists. Every endpoint causes one update, and at most m openings can be stored.

At each successful step the stack is exactly the list of chords whose opening but not closing endpoint has been read, in increasing order of opening time. This invariant holds initially for the empty list. An opening appends its label, and a successful closing deletes precisely the final still-open interval, so induction proves the invariant after every prefix. In particular, if o_j and c_j count opening and closing occurrences through position j, then

$$h_j = o_j - c_j = \#\{\text{stack entries after position } j\} \geq 0, \qquad h_{2m} = 0.$$

These unlabeled height conditions alone are insufficient. The order A, B, A, B has heights $1, 2, 1, 0$, but its third occurrence tries to close A while B is on top. The label invariant detects exactly the forbidden interleaving, whereas a balance count would accept it. Conversely, when the label test succeeds, two completed intervals are either disjoint or one contains the other: an overlap without containment would produce that same forbidden closing pattern. This provides an equivalent interval-nesting description of the successful certificate.

□

The restriction to a matching is material. At a vertex of degree three there is no unique partner occurrence to put on this stack, and an arbitrary ordering of coincident occurrences can create a false failure. The exhaustive partition test remains available for blocks of arbitrary finite size. A faster general algorithm would need its own invariant and proof, not an unannounced extension of the matching argument.

17.4 Uncertain angles and finite evidence

Suppose each true angle has a representative in a known interval $I_j \subset \mathbb{R}$ of length at most 2η, after choosing a cut outside all these intervals. If the intervals are pairwise disjoint, their order certifies the order of the true angles. The crossing decision is then valid for every choice of one true angle in each interval. If neighboring intervals overlap, an exact order conclusion may be unavailable. Replacing an overlap by an arbitrary tie-breaking rule would change the mathematical question, rather than answer it.

For example, if approximate angles differ by σ and each has absolute error at most η, their order is certified when $\sigma > 2\eta$. At equality the intervals may touch, so strict distinctness is not guaranteed. Equality of exact rational angles is decidable by (17.1); equality of unspecified real numbers represented only by successively narrower intervals need not be settled at any finite stage. One may correctly report unresolved order in the latter model.

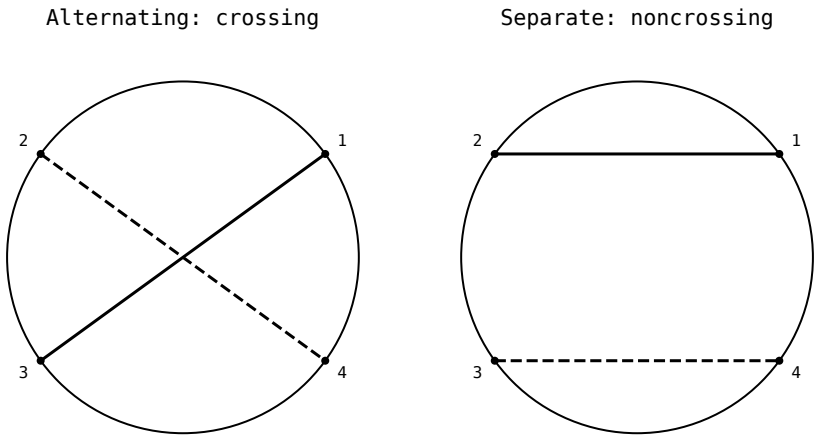


Figure 17.1: Finite chord data are decided by alternating endpoint order. Shared endpoints and repeated chords must be distinguished before applying the four-distinct-endpoint test.

For ordered approximate representatives $\widehat{t}_i < \widehat{t}_j$, the entire interval proof is the chain

$$t_j - t_i \geq (\widehat{t}_j - \eta) - (\widehat{t}_i + \eta) = \widehat{t}_j - \widehat{t}_i - 2\eta.$$

A positive right-hand side certifies both distinctness and orientation. For a finite family, take the smallest adjacent approximate gap in the chosen chart; positivity after subtracting 2η verifies all order comparisons by transitivity. The chart cut must separately lie outside every uncertainty interval, including the first and last intervals. If exact distinct real angles are accompanied by convergent enclosing intervals, their positive minimum circular separation eventually exceeds twice the interval widths. That proves eventual success for that fixed finite distinct input. It supplies no uniform precision bound over inputs whose minimum separation can approach zero, and it gives no finite certificate of equality when two angles actually coincide.

Passing every test on a finite prefix also does not certify an infinite lamination. In an unused open arc one can add four successive new endpoints and pair the first with the third and the second with the fourth. The new chords cross although the entire original finite record remains unchanged. To infer a statement about a closure of infinitely many chords requires additional hypotheses controlling all omitted pairs and limit leaves. The closed-relation and quotient arguments elsewhere in this volume address that genuinely different question.

17.5 Solved investigations

Investigation 17.5 (Crossing across the angular cut)**.** Test the chord pairs $(9/10, 1/10)$ and $(0, 1/2)$ without evaluating their Cartesian endpoints.

Solution. The counterclockwise arc from $9/10$ to $1/10$ consists of the angles greater than $9/10$ or less than $1/10$. It contains 0 and does not contain $1/2$. The Boolean values differ, so the two open chords cross. The apparent reversal $9/10 > 1/10$ merely invokes the wrapped case of J; it does not reverse the chosen orientation of the circle.

Investigation 17.6 (A shared endpoint is not a transverse crossing)**.** Compare the chords $(0, 1/3)$ and $(0, 2/3)$, and then compare $(0, 1/3)$ with $(2/6, 0)$. State the required output in each case.

Solution. The first pair shares exactly one endpoint and has disjoint interiors by Proposition 17.1. The second pair has identical unordered endpoint sets because $2/6 = 1/3$, checked by $2 \cdot 3 = 1 \cdot 6$.

It therefore represents the same geometric chord twice. The output is duplicate chord, not transverse crossing. If these were asserted to be edges belonging to disjoint partition blocks, the common endpoint would already violate the partition validation step.

Investigation 17.7 (Reading a matching certificate)**.** The endpoints of three chords appear as A, B, B, C, C, A. Compare this order with A, B, C, B, A, C.

Solution. For the first order the stack successively reads (A), (A, B), (A), (A, C), (A), and the empty stack. Every close removes the current last opening, so the matching is noncrossing. For the second order the fourth endpoint closes B while C is still on top. The witness is $B_{\text{open}} < C_{\text{open}} < B_{\text{close}} < C_{\text{close}}$, at positions $2, 3, 4, 6$. Checking only this quadruple independently confirms the rejection.

Investigation 17.8 (A partition certificate is stronger than a pair list)**.** On three indices declare 1 compatible with 2 and 2 compatible with 3, but 1 incompatible with 3. Can these pairwise records be used as blocks of a partition without further modification?

Solution. No. If 1 and 2 lie in one block and 2 and 3 lie in one block, disjointness of blocks forces those blocks to coincide, so 1 and 3 also lie together. The proposed relation is not transitive. Taking its transitive closure would give the single block $\{1, 2, 3\}$, but that is a changed relation whose geometric meaning needs separate justification. This is the finite reason that pairwise non-separation cannot silently be promoted to an equivalence relation.

Investigation 17.9 (Precision sufficient for order)**.** Approximate endpoints on a chosen linear angular chart are separated by at least $1/1000$. Each true endpoint is within 2^{-p} of its approximation. Give a sufficient integer p for all comparisons.

Solution. It suffices that $2 \cdot 2^{-p} < 1/1000$, or $2^p > 2000$. Thus $p = 11$ suffices, because $2048 > 2000$. This calculation assumes that the same chart avoids all uncertainty intervals at its cut. If an interval crosses the cut, it must first be represented as a circular interval or the cut must be moved. The bound certifies strict ordering only; it does not certify that the measured angles are the accesses of the intended continuum.

Chapter 18

Dirichlet problems and elimination on finite trees

A tree of puzzle pieces describes containment. To turn it into an electrical or thermal network one must additionally specify its terminals, edge conductances and interpretation of a potential. None of these is determined solely by the fact that the tree came from a planar drawing. We first solve the finite mathematical problem, without making that extra identification. The classical Dirichlet and current principles are discussed in Doyle and Snell [4]; the proofs here use explicit finite quadratic forms and require no limiting electrical medium.

18.1 Energy and the unique harmonic extension

Let $G = (V, E)$ be a finite connected undirected graph without loops. Give each unordered edge $\{x, y\}$ a conductance $c_{xy} = c_{yx} > 0$. Parallel edges, if present in an initial description, may be replaced by one edge with their summed conductance. Choose a nonempty boundary set $B \subset V$, write $I = V \setminus B$, and prescribe real values $g : B \to \mathbb{R}$. For $u : V \to \mathbb{R}$ define

$$\mathcal{E}_c(u) = \sum_{\{x,y\}\in E} c_{xy}(u_x - u_y)^2, \qquad (Lu)_x = \sum_{y:\{x,y\}\in E} c_{xy}(u_x - u_y). \tag{18.1}$$

Every edge is counted once in the energy. Consequently

$$\sum_{x\in V} v_x(Lu)_x = \sum_{\{x,y\}\in E} c_{xy}(v_x - v_y)(u_x - u_y), \qquad \mathcal{E}_c(u) = u^{\mathsf{T}} L u. \tag{18.2}$$

The identity follows by collecting the two endpoint terms of each edge; it is the finite summation-by-parts formula behind all subsequent proofs.

Theorem 18.1 (Finite Dirichlet principle)**.** *There is a unique function h with $h|_B = g$ and $(Lh)_x = 0$ for $x \in I$. It uniquely minimizes $\mathcal{E}_c$ among all functions with that boundary data. Every admissible u satisfies the exact energy decomposition*

$$\mathcal{E}_c(u) = \mathcal{E}_c(h) + \mathcal{E}_c(u - h).$$

If I is nonempty, its values solve the positive definite system

$$L_{II} h_I = -L_{IB} g.$$

Proof. For a vector w which vanishes on B, the energy is nonnegative and can be zero only when $w_x = w_y$ on every edge. Connectedness then makes w constant on V, and the nonempty

boundary makes this constant zero. Thus the quadratic form $w_I^{\mathsf{T}} L_{II} w_I$ is positive definite, so the displayed finite linear system has a unique solution. Extending it by g gives the asserted harmonic function. If I is empty there are no unknowns and the assertion is immediate.

For an admissible u put $w = u - h$. Expanding edge by edge gives

$$\mathcal{E}_c(h + w) = \mathcal{E}_c(h) + 2 \sum_{x \in V} w_x (Lh)_x + \mathcal{E}_c(w).$$

The middle sum vanishes: $w = 0$ on B and $Lh = 0$ on I. The remaining energy is nonnegative and is zero only when $w = 0$, by the preceding argument. This proves both minimality and uniqueness.

To expose the linear algebra behind the quadratic argument, orient each edge once and let D be the matrix with a 1 at its initial vertex and a -1 at its final vertex. With $W = \operatorname{diag}(c_e)$,

$$Du = (u_{e^-} - u_{e^+})_{e \in E}, \qquad L = D^{\mathsf{T}} W D, \qquad L_{II} = D_I^{\mathsf{T}} W D_I.$$

The last equality uses the columns indexed by I; boundary coordinates of a zero-boundary test vector are absent. If $D_I w_I = 0$, its zero-boundary extension has equal values along every edge. A path from any vertex to the nonempty boundary forces that value to be zero. Hence D_I has trivial kernel, and positivity of W gives the claimed positive definiteness without relying on a computed determinant. The energy at a general interior vector x can now be written

$$\mathcal{E}_c(x, g) = x^{\mathsf{T}} L_{II} x + 2 x^{\mathsf{T}} L_{IB} g + g^{\mathsf{T}} L_{BB} g.$$

Its gradient with respect to x is $2(L_{II} x + L_{IB} g)$. Setting this to zero gives the same harmonic equations; completing the square about their unique solution gives the same nonnegative energy gap. Thus the variational and linear-system descriptions agree for every prescribed boundary vector, including constant boundary data.

□

Proposition 18.2 (Maximum principle and stability of boundary values)**.** *The harmonic extension satisfies* $\min_B g \le h_x \le \max_B g$ *for every* $x \in V$*. For two boundary functions* $g, \widetilde{g}$ *on the same network,*

$$\max_{x \in V} |h_x - \widetilde{h}_x| \le \max_{b \in B} |g_b - \widetilde{g}_b|.$$

Proof. At an interior vertex the equation $Lh = 0$ can be solved as

$$h_x = \sum_{y \sim x} \frac{c_{xy}}{\sum_{z \sim x} c_{xz}} h_y.$$

All weights are positive and sum to one. If a global maximum strictly larger than all boundary values were attained at an interior vertex, all its neighbors would have the same maximum. Propagating along a path to B gives a contradiction. Apply the same reasoning to $-h$ for the minimum. The difference $h - \widetilde{h}$ is harmonic in I by linearity, so applying both bounds to its boundary data proves the stability estimate.

The propagation step can be checked by subtracting each neighbor value from a maximal value $M = h_x$:

$$0 = \sum_{y \sim x} c_{xy} (M - h_y), \qquad c_{xy} > 0, \quad M - h_y \ge 0.$$

Every summand must vanish, so each neighbor also has value M. Connectedness supplies a finite path to the boundary, and induction along its successive edges reaches the contradiction in finitely many steps. Zero conductances would invalidate that propagation along an edge, which is why the stated positivity and connectedness matter. There is also a useful matrix reading.

Extend each boundary coordinate vector harmonically and write its value at x as H_{xb}. The just-proved bounds and linearity imply

$$H_{xb} \geq 0, \qquad \sum_{b \in B} H_{xb} = 1, \qquad h_x = \sum_{b \in B} H_{xb} g_b.$$

The sum is one because the sum of boundary coordinate vectors is the constant-one boundary vector. Thus every interior value is a convex combination of boundary values, and the stability constant one follows by applying the triangle inequality to this explicit representation. This finite kernel requires no probabilistic interpretation.

□

18.2 Eliminating one vertex exactly

Suppose an interior vertex has neighboring potentials $v_1, \ldots, v_k$ and incident conductances $c_1, \ldots, c_k$. Its own potential x appears in the local energy $\sum_j c_j (x - v_j)^2$. Writing $C = \sum_j c_j > 0$ and $m = C^{-1} \sum_j c_j v_j$ gives

$$\sum_{j=1}^{k} c_j (x - v_j)^2 = C(x - m)^2 + \sum_j c_j v_j^2 - C m^2, \tag{18.3}$$

$$\min_x \sum_{j=1}^{k} c_j (x - v_j)^2 = \frac{1}{C} \sum_{1 \leq i < j \leq k} c_i c_j (v_i - v_j)^2. \tag{18.4}$$

For completeness, multiplying the second constant term by C and expanding produces

$$C \sum_j c_j v_j^2 - \left(\sum_j c_j v_j \right)^2 = \sum_{i<j} c_i c_j (v_i^2 + v_j^2 - 2 v_i v_j),$$

which proves the identity. Thus eliminating a vertex creates conductance $c_i c_j / C$ between each pair of its neighbors. Existing conductances between the same pair are added. Even if the original graph is a tree, eliminating a branching vertex can create a triangle, so the reduced graph need not remain a tree.

For $k = 2$ the new conductance is

$$c_{\text{series}} = \frac{c_1 c_2}{c_1 + c_2}, \qquad \frac{1}{c_{\text{series}}} = \frac{1}{c_1} + \frac{1}{c_2}. \tag{18.5}$$

For $k = 1$ the minimized local energy is zero: a free leaf takes the potential of its only neighbor and draws no current. A leaf with a prescribed value is different; its incident energy must be retained, since its potential is not a variable over which one may minimize.

Proposition 18.3 (The boundary response matrix)**.** *If $I \neq \varnothing$, define*

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

If $I = \varnothing$, put $\Lambda = L$. The vector of currents flowing from boundary vertices into the network is Λg, and the minimum energy is

$$\mathcal{E}_c(h) = g^{\mathsf{T}} \Lambda g.$$

The matrix Λ is symmetric and positive semidefinite, its row sums are zero, and its off-diagonal entries are nonpositive.

Proof. Substitute $h_I = -L_{II}^{-1}L_{IB}g$ into the boundary rows of Lh to obtain $(Lh)_B = \Lambda g$. Formula (18.2) and $Lh = 0$ on I then give the energy identity. Symmetry follows from that of L and L_{II}^{-1}, and nonnegativity from the energy. Constant boundary data extend constantly, so $\Lambda\mathbf{1} = 0$. For distinct boundary vertices b, j, set $g_j = 1$ and all other boundary values to zero. The maximum principle gives $0 \le h \le 1$, and therefore

$$\Lambda_{bj} = (Lh)_b = \sum_{y\sim b} c_{by}(0 - h_y) \le 0.$$

This proves the sign assertion, including any direct boundary edges. □

Thus internal vertices may be removed while preserving all boundary voltage-current relations, rather than merely one terminal measurement. When all conductances and prescribed values are rational, the equations and elimination identities also show that every resulting potential and response coefficient is rational. This is an exact algebraic assertion; it does not bound the growth of numerators and denominators in a particular implementation of elimination.

18.2.1 Two interior vertices: a complete rational elimination

An exact small network illustrates all boundary entries at once. Join terminals a, b to an interior vertex u, terminals c, d to an interior vertex v, and join u to v. Give all five edges conductance one. In the orders (u, v) and (a, b, c, d), the blocks are

$$L_{II} = \begin{pmatrix} 3 & -1 \\ -1 & 3 \end{pmatrix}, \qquad L_{II}^{-1} = \frac{1}{8}\begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}, \qquad L_{IB} = -\begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix}, \quad L_{BB} = I_4.$$

Multiplication verifies the inverse because the diagonal products are $9-1 = 8$ and the off-diagonal products are $3 - 3 = 0$. Therefore

$$\begin{pmatrix} u \\ v \end{pmatrix} = \frac{1}{8}\begin{pmatrix} 3(g_a + g_b) + (g_c + g_d) \\ (g_a + g_b) + 3(g_c + g_d) \end{pmatrix},$$

and direct block multiplication gives the full response

$$\Lambda = \frac{1}{8}\begin{pmatrix} 5 & -3 & -1 & -1 \\ -3 & 5 & -1 & -1 \\ -1 & -1 & 5 & -3 \\ -1 & -1 & -3 & 5 \end{pmatrix}.$$

Every row sums to zero. Reading the negative off-diagonal entries as conductances gives $3/8$ on ab and cd, and $1/8$ on each of the four pairs with one endpoint in $\{a, b\}$ and the other in $\{c, d\}$. Successive star elimination gives an independent check. Eliminating u first creates conductance $1/3$ on ab, av, bv. The remaining incident conductance at v is $1/3 + 1/3 + 1 + 1 = 8/3$, so its elimination gives

$$c_{ab} = \frac{1}{3} + \frac{(1/3)^2}{8/3} = \frac{3}{8}, \qquad c_{ac} = \frac{(1/3)\cdot 1}{8/3} = \frac{1}{8}, \qquad c_{cd} = \frac{1\cdot 1}{8/3} = \frac{3}{8}.$$

The other cross-pairs follow by the displayed symmetries. For terminal values $(0, 1, 2, 3)$, the recovered interiors are $(u, v) = (1, 2)$ and

$$\Lambda(0, 1, 2, 3)^{\mathsf{T}} = (-1, 0, 0, 1)^{\mathsf{T}}, \qquad \mathcal{E} = 1 + 0 + 1 + 0 + 1 = 3.$$

The response energy also equals $(0, 1, 2, 3) \cdot (-1, 0, 0, 1) = 3$. This agreement holds for every boundary vector by the matrix identity, not just for these four chosen voltages. The reduced graph generally has different internal topology; equality of terminal response does not make its drawing a conformal or metric image of the original tree.

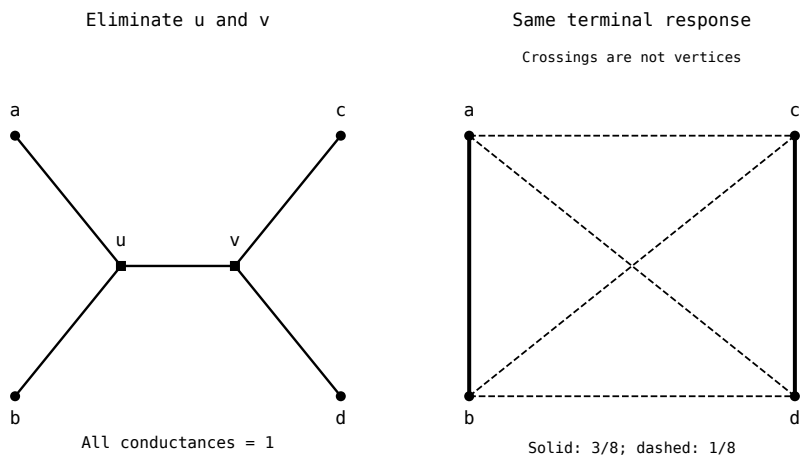


Figure 18.1: The five unit-conductance edges of an H-shaped tree reduce to terminal conductances $3/8$ on ab, cd and $1/8$ on the four cross-pairs. This is an exact boundary-response identity, not a geometric identification of the drawings.

18.3 Bottom-up messages on a rooted tree

A rooted tree permits elimination in an order which keeps only a quadratic function at each active parent. For a subtree below a vertex with potential x, write its already minimized energy as

$$Q(x) = Ax^2 - 2Bx + C.$$

This function is nonnegative for every real x, because it arose by minimizing a sum of nonnegative edge energies. Hence $A \geq 0$, and if $A = 0$ then $B = 0$. Connect this vertex to its parent, of potential y, through conductance $c > 0$. Completing the square again gives

$$\min_x \{c(y-x)^2 + Q(x)\} = A'y^2 - 2B'y + C', \tag{18.6}$$

$$A' = \frac{cA}{c+A}, \qquad B' = \frac{cB}{c+A}, \qquad C' = C - \frac{B^2}{c+A}, \tag{18.7}$$

$$x_{\min} = \frac{cy+B}{c+A}. \tag{18.8}$$

Indeed, the quadratic coefficient of x is $c + A > 0$, its linear coefficient is $-2(cy + B)$, and its remaining term is $cy^2 + C$. Subtracting $(cy + B)^2/(c + A)$ yields the coefficients above.

For $A > 0$, rewrite the incoming message in a form which isolates its preferred potential and unavoidable constant cost:

$$Q(x) = A(x-m)^2 + D_0, \qquad m = \frac{B}{A}, \qquad D_0 = C - \frac{B^2}{A} \geq 0.$$

Then the complete square is

$$c(y-x)^2 + A(x-m)^2 = (c+A)\left(x - \frac{cy+Am}{c+A}\right)^2 + \frac{cA}{c+A}(y-m)^2.$$

Thus transmission preserves the preferred value m and constant cost D_0, while changing the positive stiffness from A to $cA/(c + A)$. In coefficient form the same fact is

$$\frac{B'}{A'} = \frac{B}{A}, \qquad C' - \frac{(B')^2}{A'} = C - \frac{B^2}{A}.$$

If $A = 0$, nonnegativity of Q on all of $\mathbb{R}$ forces $B = 0$, so the message is constant and the minimizer is $x = y$. This separate case avoids dividing by A. At a branching vertex the incoming constants and stiffnesses add before transmission; different preferred values can then leave a strictly positive minimized constant cost. Back substitution records the weighted compromise at that vertex rather than pretending that every branch preferred the same voltage.

A prescribed leaf value g sends the message $c(y - g)^2$, with coefficients (c, cg, cg^2). A free leaf sends zero. At a vertex, the messages of its children add coefficient by coefficient; then (18.7) sends the combined message to its parent. With fixed data on the leaves and, if desired, the root, this procedure eliminates every unprescribed vertex. If the root is free and its resulting quadratic has $A > 0$, its value is B/A. When at least one leaf is fixed, connectedness and positive conductances force that final A to be positive: otherwise the minimum energy would be independent of the root value, in contradiction with the uniqueness of the Dirichlet extension for varying root candidates. Equivalently, the recursive positive coefficients along a path to a fixed leaf remain positive.

Back substitution using (18.8) recovers every vertex potential. Each edge participates in one upward message and one downward substitution, giving $O(|V|)$ arithmetic operations for this tree procedure. This count treats a rational arithmetic operation as one operation; exact bit complexity additionally depends on the sizes of its operands. Prescribed interior vertices can be treated as separators, solving the components between them with those values fixed.

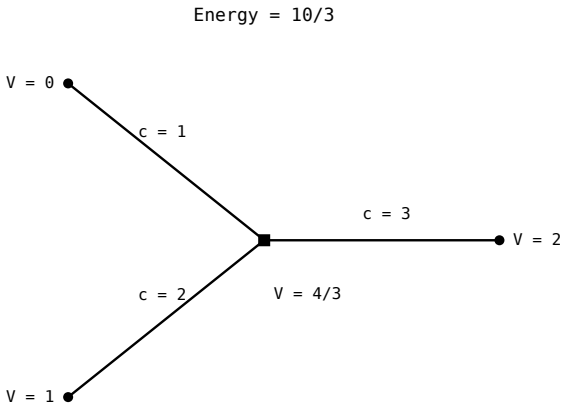


Figure 18.2: A three-arm network with conductances $1, 2, 3$ and terminal potentials $0, 1, 2$. The central potential is $4/3$ and the once-per-edge energy is $10/3$ in the chosen normalized units.

18.4 Solved investigations

Investigation 18.4 (The three-arm star)**.** Three terminal values are $0, 1, 2$, with incident conductances $1, 2, 3$. Find the central potential, each outward terminal current and the energy.

Solution. The center is the weighted mean $x = (1 \cdot 0 + 2 \cdot 1 + 3 \cdot 2)/(1 + 2 + 3) = 4/3$. Currents from the terminals into the network are $1(0 - x) = -4/3$, $2(1 - x) = -2/3$, and $3(2 - x) = 2$. Their sum is zero. The energy is

$$(4/3)^2 + 2(1/3)^2 + 3(2/3)^2 = 10/3.$$

The boundary power identity gives the same value: $0(-4/3) + 1(-2/3) + 2(2) = 10/3$. This second calculation checks the current signs and edge-count convention.

Investigation 18.5 (A series path with a free branch)**.** A path from terminal a to terminal b has edge resistances $1, 2, 3$. A finite tree with no additional terminal is attached to its middle vertex. Compute the effective conductance between a and b.

Solution. Successively eliminate free leaves in the attached tree. Every eliminated leaf has zero minimized energy, so eventually the entire added tree disappears without changing the boundary response. The path resistance is $1 + 2 + 3 = 6$ by (18.5), hence its conductance is $1/6$. For a unit voltage the current is $1/6$, and the drops across the three edges are $1/6, 2/6, 3/6$. Every attached branch takes the potential of its attachment vertex. A new prescribed terminal on such a branch would change this calculation and cannot be discarded.

Investigation 18.6 (A rational message)**.** A child subtree sends $Q(x) = 2x^2 - 6x + 5$ to a vertex linked to its parent by conductance 3. Compute its message and verify nonnegativity.

Solution. Here $(A, B, C) = (2, 3, 5)$, so

$$Q_{\text{parent}}(y) = \frac{6}{5}y^2 - \frac{18}{5}y + \frac{16}{5} = \frac{6}{5}\left(y - \frac{3}{2}\right)^2 + \frac{1}{2}.$$

The original quadratic is $2(x - 3/2)^2 + 1/2$, also nonnegative. At a fixed parent value y, the child potential is $(3y + 3)/5$. Substitution into $3(y - x)^2 + Q(x)$ recovers the displayed message, confirming that its constant term has not been lost during elimination.

Investigation 18.7 (The response of a star)**.** For a center joined to k boundary vertices with conductances c_j, write the complete response matrix and identify its nullspace.

Solution. Put $c = (c_1, \dots, c_k)^{\mathsf{T}}$ and $C = \sum_j c_j$. Elimination gives $\Lambda = \operatorname{diag}(c_1, \dots, c_k) - cc^{\mathsf{T}}/C$. Its energy is $C^{-1} \sum_{i<j} c_i c_j (g_i - g_j)^2$. For $k \geq 2$, positivity of every coefficient makes this zero exactly when all g_i coincide; the nullspace is the constants. For $k = 1$ the matrix is the one-by-one zero matrix, whose entire one-dimensional space is again the space of constants. Each off-diagonal entry is $-c_i c_j / C$, consistent with Proposition 18.3.

Investigation 18.8 (Why anchoring matters)**.** Remove all prescribed boundary values from a connected network. Describe the minimizers of its energy and explain why the previous uniqueness statement no longer applies.

Solution. Every constant function has energy zero, and zero energy forces equality along every edge, hence constancy by connectedness. Thus there is a whole one-parameter family of minimizers. The full Laplacian has $L\mathbf{1} = 0$ and is not invertible. The nonempty boundary in Theorem 18.1 removed this additive freedom. Prescribing one value restores uniqueness, though it may produce the constant solution if there are no other imposed data or sources.

Chapter 19

Resistance bounds, perturbations and finite error certificates

The preceding elimination identities compute a finite network exactly when rational arithmetic is feasible. Bounds remain useful when one does not want to eliminate the whole graph, when conductances are uncertain, or when an approximate potential has already been proposed. In each case an acceptance argument must be independent of the procedure that generated the proposal. The central quantities are variational, which makes such arguments short enough to inspect directly.

19.1 Voltage and current formulations

Choose distinct terminals $a, b \in V$, prescribe $h_a = 1$, $h_b = 0$, and let all other vertices be harmonic. Define

$$C_{ab} = \mathcal{E}_c(h), \qquad R_{ab} = C_{ab}^{-1}.$$

Connectedness and distinct terminal values imply $C_{ab} > 0$. By (18.2), $C_{ab} = (Lh)_a$, the current entering at a under unit voltage. The current at b is $-C_{ab}$ because the sum of all rows of Lh is zero. Thus the reciprocal definition agrees with voltage divided by current.

Temporarily choose one orientation for every edge and write e^-, e^+ for its initial and final vertices. A signed flow is a vector (j_e); its divergence is

$$(\operatorname{div} j)_x = \sum_{e:e^-=x} j_e - \sum_{e:e^+=x} j_e.$$

A unit flow from a to b has divergence 1 at a, -1 at b, and zero elsewhere. With $r_e = 1/c_e$, its dissipation is $\mathcal{D}_r(j) = \sum_e r_e j_e^2$. Changing an edge orientation changes the sign of its flow and leaves this quantity invariant. Collecting endpoint terms gives a second useful summation-by-parts formula:

$$\sum_e (u_{e^-} - u_{e^+}) j_e = \sum_x u_x (\operatorname{div} j)_x. \tag{19.1}$$

Theorem 19.1 (Unit-flow principle)**.** *Among all unit flows from a to b, the unique minimizer of $\mathcal{D}_r$ is*

$$i_e = \frac{c_e(h_{e^-} - h_{e^+})}{C_{ab}}, \qquad \min_j \mathcal{D}_r(j) = R_{ab}.$$

For any unit flow j, the exact orthogonal decomposition is

$$\mathcal{D}_r(j) = R_{ab} + \mathcal{D}_r(j - i).$$

Proof. The divergence equations for h show that i is a unit flow. Its dissipation is

$$\sum_e \frac{1}{c_e}\frac{c_e^2(h_{e^-}-h_{e^+})^2}{C_{ab}^2} = \frac{\mathcal{E}_c(h)}{C_{ab}^2} = \frac{1}{C_{ab}}.$$

For another unit flow put $k = j - i$; then $\operatorname{div} k = 0$ at every vertex, including the terminals. The cross term in the expansion of $\mathcal{D}_r(i+k)$ is

$$\sum_e r_e i_e k_e = \frac{1}{C_{ab}}\sum_e (h_{e^-}-h_{e^+})k_e = \frac{1}{C_{ab}}\sum_x h_x(\operatorname{div} k)_x = 0.$$

The expansion proves the identity. All resistances are positive, so its last term vanishes only if $k = 0$, proving uniqueness. □

The voltage formulation supplies an upper bound on C_{ab} from any potential with the correct terminal values. The current formulation supplies an upper bound on R_{ab} from any unit flow. Accordingly a trial potential v and a trial unit flow j yield the two-sided enclosure

$$\frac{1}{\mathcal{E}_c(v)} \le R_{ab} \le \mathcal{D}_r(j). \tag{19.2}$$

The left denominator is positive because the endpoints differ. These are exact inequalities if the feasibility conditions and energies have been checked exactly or enclosed with justified directed errors. A small mismatch in divergence is not a unit-flow certificate until it has been corrected or its effect bounded.

The two trial certificates also satisfy a direct product inequality. Since $v_a - v_b = 1$ and $\operatorname{div} j = \mathbf{1}_a - \mathbf{1}_b$, summation by parts and weighted Cauchy–Schwarz give

$$1 = \sum_e \Delta_e v\, j_e, \qquad 1 \le \left(\sum_e c_e(\Delta_e v)^2\right)\left(\sum_e r_e j_e^2\right) = \mathcal{E}_c(v)\mathcal{D}_r(j).$$

Here $\Delta_e v = v_{e^-} - v_{e^+}$, with the same orientation as j_e. Equality in Cauchy–Schwarz requires j_e to be one common multiple of $c_e\Delta_e v$ on all edges. Together with feasibility, this means that v is the harmonic unit-voltage potential and j its normalized unit flow. More quantitatively, set $E = \mathcal{E}_c(v)$ and $D = \mathcal{D}_r(j)$. The interval width decomposes into two nonnegative defects:

$$D - \frac{1}{E} = (D - R_{ab}) + \frac{E - C_{ab}}{C_{ab}E}.$$

The first defect is exactly the flow energy of $j - i$, and the numerator in the second is exactly the potential energy of $v - h$. This explains how independent feasible trials can squeeze the resistance from both sides. It also shows why raw residuals without feasibility do not give the same enclosure: the initial summation-by-parts value would cease to equal one.

Proposition 19.2 (Paths and cuts)**.** *For any simple path P from a to b, and any vertex set S containing a but not b, let ∂S be the edges with one endpoint in S. Then*

$$\frac{1}{\sum_{e\in\partial S} c_e} \le R_{ab} \le \sum_{e\in P} r_e.$$

On a tree the upper bound is an equality for its unique terminal path.

Proof. Send one unit of flow along P and zero elsewhere. This is feasible and has dissipation $\sum_{e\in P} r_e$, giving the upper bound. For the lower bound take the trial potential $v = 1$ on S and $v = 0$ outside. Its energy is exactly $\sum_{e\in\partial S} c_e$, so (19.2) applies. On a tree, deleting an edge divides the vertices into two components. Summing divergence over either component shows that every unit flow has current of magnitude one through the edge if it separates a from b, and zero otherwise. The separating edges are precisely those of the unique path. Thus its path flow is the only unit flow, proving equality. □

This also derives the series law without eliminating potentials, and explains why a free dangling tree cannot affect a two-terminal resistance. The argument uses the absence of a cycle. A graph with an additional path can divide its unit flow and lower the resistance.

19.1.1 Several separating cuts give an additive lower bound

Proposition 19.3 (Disjoint cut bounds)**.** *Let $S_1, \dots, S_k$ contain a and exclude b. Suppose their edge cuts $\partial S_1, \dots, \partial S_k$ are pairwise disjoint as sets of edges. Then*

$$R_{ab} \geq \sum_{j=1}^{k} \frac{1}{\sum_{e \in \partial S_j} c_e}.$$

The sets S_j need not be disjoint; disjointness is required of their crossing edges.

Proof. For a unit flow i, orient each edge of a selected cut outward from S_j for the purpose of that cut's calculation. Summing divergence over S_j cancels every internal edge and gives $\sum_{e \in \partial S_j} i_e = 1$. Weighted Cauchy–Schwarz gives

$$1 = \left(\sum_{e \in \partial S_j} \sqrt{r_e} i_e \sqrt{c_e} \right)^2 \leq \left(\sum_{e \in \partial S_j} r_e i_e^2 \right) \left(\sum_{e \in \partial S_j} c_e \right).$$

Because the cuts share no edges, their dissipation sums add without counting any edge twice, and their total is at most $\mathcal{D}_r(i)$. Hence every unit flow has dissipation at least the displayed sum of reciprocal cut conductances. Minimize over flows using Theorem 19.1. □

For a path with successive resistances $r_1, \dots, r_k$, take cuts at its successive edges. The lower bound becomes $\sum_j r_j$, matching the path-flow upper bound. For a branching graph, broad cuts can have large total conductance and hence contribute only a small reciprocal term. This is the quantitative distinction between several bottlenecks in series and many alternative edges in one bottleneck.

The disjointness hypothesis cannot be removed while retaining the sum. On a single edge of resistance one, using the same cut twice would incorrectly predict resistance at least two. More subtly, nested vertex sets can still have overlapping edge cuts when an edge jumps across several layers. Nestedness of regions alone does not justify adding their cut bounds. A finite certificate can list the cut edges and check that each edge index appears at most once; this is an exact combinatorial condition separate from estimating the positive conductance sums.

19.2 Monotonicity and uncertainty in conductances

Proposition 19.4 (Rayleigh comparison and relative perturbations)**.** *On a fixed connected graph with fixed distinct terminals, increasing any conductances cannot decrease C_{ab} and cannot increase R_{ab}. If, for $0 \leq \eta < 1$,*

$$(1 - \eta) c_e \leq \widetilde{c}_e \leq (1 + \eta) c_e \quad \textit{for every edge,}$$

then

$$(1 - \eta) C_{ab} \leq \widetilde{C}_{ab} \leq (1 + \eta) C_{ab},$$
$$\frac{R_{ab}}{1 + \eta} \leq \widetilde{R}_{ab} \leq \frac{R_{ab}}{1 - \eta}.$$

Proof. For each admissible potential, its energy is an increasing function of each conductance. Taking the infimum over the same set of potentials preserves this ordering. Taking positive reciprocals reverses it for resistance. Under the relative hypotheses, every admissible potential satisfies

$$(1-\eta)\mathcal{E}_c(u) \le \mathcal{E}_{\widetilde{c}}(u) \le (1+\eta)\mathcal{E}_c(u).$$

Minimizing yields the conductance bounds, and reciprocation yields the resistance bounds. This proof does not assume that the two minimizing potentials are equal. □

The same argument shows that C_{ab} is concave as a function of the conductance vector: for $0 \le t \le 1$,

$$\inf_u \{t\mathcal{E}_c(u) + (1-t)\mathcal{E}_d(u)\} \ge t \inf_u \mathcal{E}_c(u) + (1-t)\inf_u \mathcal{E}_d(u).$$

The unit-flow formula similarly makes R_{ab} concave in the resistance vector. These are different parameterizations; concavity in conductance must not be silently replaced by concavity in resistance for the same output quantity.

There is a precise local sensitivity formula. If one conductance c_e varies over positive values while the others are fixed, the positive definite linear system makes its harmonic solution differentiable. Differentiating the minimized energy gives

$$\frac{\partial C_{ab}}{\partial c_e} = (h_{e^-} - h_{e^+})^2. \tag{19.3}$$

Indeed, the derivative from changing c_e at fixed h is the displayed square. The terms from differentiating h sum to $2\sum_x h'_x(Lh)_x$, which vanishes because h' is zero at the terminals and Lh is zero elsewhere. The differentiability assertion can also be read directly from the inverse matrix formula: its entries are quotients of polynomials with a nonzero determinant on positive conductances. Thus an edge with zero potential drop has zero first-order effect, although its effect at other conductance values must still be considered.

Differentiating the interior equations makes the dependence precise. Along a differentiable positive-conductance path $c(t)$ with fixed terminal values, write a dot for the derivative. Then

$$L_{II}\dot{h}_I = -(\dot{L}h)_I, \qquad \dot{h}_a = \dot{h}_b = 0,$$

so the derivative is determined by the same invertible grounded matrix. In the differentiated energy the solution-motion contribution is $2\dot{h}^{\mathsf{T}}Lh = 0$, leaving

$$\frac{\mathrm{d}C_{ab}}{\mathrm{d}t} = \sum_e \dot{c}_e(\Delta_e h)^2, \qquad \sum_e c_e \frac{\partial C_{ab}}{\partial c_e} = C_{ab}.$$

The second identity follows either by setting $\dot{c} = c$ or directly from the energy formula. Hence the dimensionless sensitivities $s_e = c_e(\Delta_e h)^2/C_{ab}$ are nonnegative and sum to one. They describe the local distribution of conductance sensitivity, not probabilities of a physical event. For resistance parameters, the chain rule and $c_e = 1/r_e$ yield

$$\frac{\partial R_{ab}}{\partial r_e} = \frac{(c_e\Delta_e h)^2}{C_{ab}^2} = i_e^2, \qquad \sum_e r_e \frac{\partial R_{ab}}{\partial r_e} = R_{ab}.$$

The derivatives are local statements at positive parameters. Finite changes require integration along a positive path or the global Rayleigh comparison already proved; multiplying a local derivative by a large parameter change without controlling its variation is not such a proof.

19.3 A residual bound with an explicit stability constant

Let B again be any nonempty boundary set and assume $I \neq \varnothing$. For each $x \in I$ choose a simple path P_x to some boundary vertex and set

$$m = |I|, \qquad R_* = \max_{x \in I} \sum_{e \in P_x} r_e.$$

All these quantities concern the supplied finite graph. In particular R_* is available from explicit paths; it is not a hidden asymptotic constant. For any w vanishing on B, telescoping along P_x and Cauchy–Schwarz give

$$|w_x|^2 \leq \left(\sum_{e \in P_x} r_e \right) \left(\sum_{e \in P_x} c_e (\Delta_e w)^2 \right) \leq R_* \mathcal{E}_c(w), \tag{19.4}$$

$$\|w_I\|_2^2 \leq m R_* \mathcal{E}_c(w). \tag{19.5}$$

In particular, the smallest eigenvalue of L_{II} is at least $(mR_*)^{-1}$. The path estimate is deliberately conservative and often can be sharpened, but it is fully explicit.

Proposition 19.5 (Residual-to-error conversion)**.** *Let $\widetilde{h}$ have exactly the prescribed boundary values and put $r = (L\widetilde{h})_I$. For the exact harmonic extension h,*

$$\|\widetilde{h}_I - h_I\|_2 \leq m R_* \|r\|_2, \qquad \mathcal{E}_c(\widetilde{h} - h) \leq m R_* \|r\|_2^2.$$

Moreover, $\mathcal{E}_c(\widetilde{h}) - \mathcal{E}_c(h) = \mathcal{E}_c(\widetilde{h} - h)$ exactly.

Proof. Writing $e = \widetilde{h} - h$, we have $e_B = 0$ and $L_{II} e_I = r$. Formula (19.5) and summation by parts imply

$$\frac{\|e_I\|_2^2}{m R_*} \leq e_I^{\mathsf{T}} L_{II} e_I = e_I^{\mathsf{T}} r \leq \|e_I\|_2 \|r\|_2.$$

If $e_I = 0$ the claims are immediate; otherwise division gives the first bound. Substituting it into $\mathcal{E}_c(e) = e_I^{\mathsf{T}} r$ gives the second. The exact energy-gap identity is Theorem 18.1 applied to $\widetilde{h}$.

An eigenbasis makes both error quantities transparent. Write $L_{II} v_j = \lambda_j v_j$ for an orthonormal basis, with $\lambda_j > 0$, and put $r_j = v_j^{\mathsf{T}} r$. The exact identities are

$$e_I = \sum_j \frac{r_j}{\lambda_j} v_j, \qquad \|e_I\|_2^2 = \sum_j \frac{r_j^2}{\lambda_j^2}, \qquad \mathcal{E}_c(e) = \sum_j \frac{r_j^2}{\lambda_j}.$$

The path inequality supplies $\lambda_j \geq (mR_*)^{-1}$ for every j, giving the claimed bounds term by term. In particular the energy error is not the square of the Euclidean error with an unspecified constant; it is the exact quadratic quantity $r^{\mathsf{T}} L_{II}^{-1} r$. The argument also isolates the zero-residual case: if $r = 0$, each coefficient r_j/λ_j is zero, and uniqueness gives $\widetilde{h} = h$ exactly. For nonzero residual, the orientation chosen to write edge differences is irrelevant because the energy squares those differences. The simple paths used to define R_* ensure that no edge contribution is counted twice inside any one path estimate.

□

If boundary values are also uncertain by at most η in maximum norm, first compare exact harmonic extensions using Proposition 18.2, then apply the residual bound to the approximate extension for its own boundary values. The triangle inequality yields the maximum-norm estimate

$$\|\widetilde{h} - h\|_\infty \leq \eta + m R_* \|r\|_2.$$

Here the residual must be calculated for the boundary values actually used. Floating arithmetic that merely prints a small residual is not enough: one also needs a justified enclosure for that residual and the positive conductances on which the stability estimate depends.

19.4 A complete branching model

Proposition 19.6 (Root-to-level resistance of a regular tree)**.** *Let every nonterminal vertex of a rooted tree have $b \geq 2$ children. All edges joining depth $k-1$ to depth k have resistance $r_k > 0$. At depth n, identify all leaves as one terminal, and take the root as the other. Then*

$$R_n = \sum_{k=1}^{n} \frac{r_k}{b^k}.$$

The identification is part of the electrical boundary condition; it is not an assertion that the corresponding puzzle points coincide.

Proof. For any unit flow to the leaf terminal, the signed currents across the b^k edges of level k sum to one by conservation at all vertices above that level. Cauchy–Schwarz implies

$$\sum_{e \text{ at level } k} j_e^2 \geq \frac{1}{b^k}.$$

Multiplying by r_k and summing gives dissipation at least the proposed value. The flow which assigns current b^{-k} to each outward level-k edge is feasible: every interior vertex receives $b^{-(k-1)}$ and sends $b \cdot b^{-k}$ onward. Its dissipation attains that lower bound. The unit-flow principle proves the formula. □

For constant $r_k = 1$, the result is $R_n = (1 - b^{-n})/(b-1)$. For $r_k = b^k$, it is $R_n = n$. Thus the same branching combinatorics can have bounded or unbounded resistance as depth increases. Neither outcome determines the local connectivity of a geometric realization. The edge weights carry analytic information absent from the bare address tree.

The corresponding potentials can be recovered level by level. Under unit current, every edge at level k has current b^{-k}, hence voltage drop $r_k b^{-k}$. If the wired leaves have potential zero, the common potential at depth j is

$$U_j = \sum_{k=j+1}^{n} \frac{r_k}{b^k}, \qquad U_0 = R_n, \qquad U_n = 0, \qquad U_{j-1} - U_j = \frac{r_j}{b^j}.$$

At an interior vertex the entering current equals the sum of its b outgoing currents, so these values satisfy every harmonic equation. Dividing by $R_n > 0$ converts them to unit-voltage boundary data. This is a second route from the proposed symmetric current to the exact Dirichlet solution, and verifies that wiring leaves means imposing a common terminal potential rather than deleting all their incident edges. For $b = 2$, three levels give respectively

$$R_3 = \frac{1}{2} + \frac{1}{4} + \frac{1}{8} = \frac{7}{8} \quad (r_k = 1), \qquad R_3 = 1 + 1 + 1 = 3 \quad (r_k = 2^k).$$

In the first case the omitted tail of the infinite series is exactly 2^{-n}; in the second case each additional layer adds exactly one. These are consequences of specified weights at every depth. A finite tree drawing by itself determines neither continuation.

19.5 Solved investigations

Investigation 19.7 (Independent upper and lower certificates)**.** Two terminals are joined by a direct resistance 3 and by a separate two-edge path with resistances 1 and 2. Find the exact resistance and exhibit a trial potential and flow giving matching bounds.

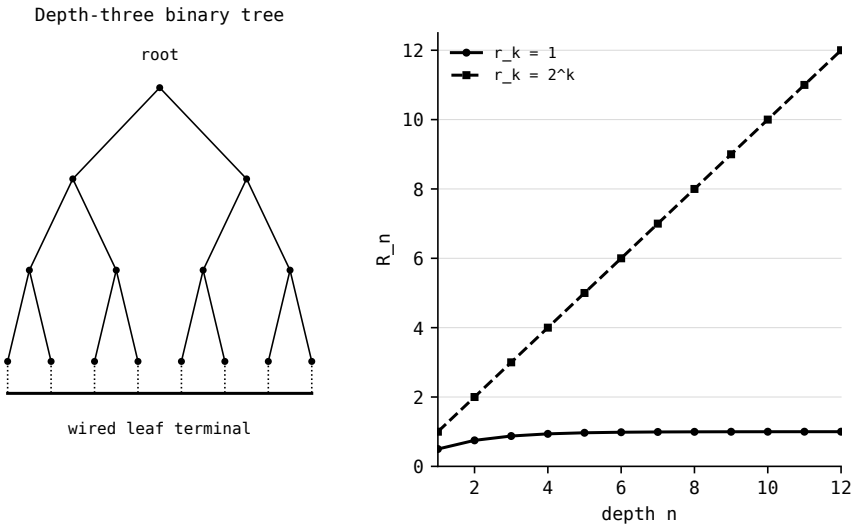


Figure 19.1: A finite binary tree with its terminal leaves wired together, and the exact level formulas $R_n = 1 - 2^{-n}$ for $r_k = 1$ and $R_n = n$ for $r_k = 2^k$. The curves concern specified weighted models, not measured resistances or a universal law for puzzle trees.

Solution. Both paths have resistance 3, so splitting a unit current equally gives dissipation $3(1/2)^2 + (1 + 2)(1/2)^2 = 3/2$. Set terminal potentials to 1 and 0 and the intermediate path potential to 2/3, placing the resistance-1 edge next to the first terminal. The energy is $1/3 + (1/3)^2 + (1/2)(2/3)^2 = 2/3$. Thus (19.2) gives $3/2 \le R_{ab} \le 3/2$. The matching witnesses prove the value without relying on the rule for parallel resistances as an unproved additional input.

Investigation 19.8 (A small residual with a large error)**.** One interior vertex is joined to a boundary vertex of value zero by conductance $\varepsilon > 0$. Test the candidate interior potential one.

Solution. The exact potential is zero, while the candidate residual is $r = \varepsilon$. Its error remains one as $\varepsilon \downarrow 0$. Here $m = 1$ and $R_* = 1/\varepsilon$, so $mR_*|r| = 1$; the bound of Proposition 19.5 is exact. The example isolates the missing stability factor in any claim that a small absolute residual alone proves a small absolute solution error.

Investigation 19.9 (Conductance uncertainty)**.** Every conductance is known to relative accuracy at most 1/20. If the nominal resistance is 7, enclose the true resistance and give an upper bound for its relative error from the nominal value.

Solution. Proposition 19.4 gives $140/21 \le \widetilde{R} \le 140/19$. Relative to 7, the upward change is at most 1/19 and the downward change is at most 1/21. Therefore the absolute relative error is at most 1/19, not automatically 1/20. Reciprocation of conductance uncertainty creates asymmetric resistance bounds.

Investigation 19.10 (Two weights on one binary tree)**.** For the complete binary tree of depth n, compute root-to-wired-leaf resistance when $r_k = 1$ and when $r_k = 2^k$. What conclusion follows about resistance from a picture of the unweighted tree?

Solution. The values are $1 - 2^{-n}$ and n, respectively, by Proposition 19.6. The same finite combinatorial tree therefore realizes sharply different values and different limiting behavior under different positive weights. A picture without weight data supplies neither value. Even if every geometric cell in a related puzzle shrinks, that fact does not choose between these two assignments.

Investigation 19.11 (An edge-sensitivity check)**.** Two equal conductances c form a series path between two terminals. Differentiate its effective conductance with respect to one conductance, holding the other fixed, and compare with (19.3).

Solution. For independent values c_1, c_2 the effective conductance is $c_1 c_2/(c_1 + c_2)$. Its derivative with respect to c_1 is $c_2^2/(c_1 + c_2)^2$, which equals $1/4$ at $c_1 = c_2 = c$. At unit terminal voltage the middle potential is $1/2$, so the squared drop on either edge is also $1/4$. If both conductances vary together, the total derivative is $1/2$, agreeing with $C(c, c) = c/2$. Distinguishing a partial derivative from a simultaneous variation is essential in using a sensitivity formula.

Chapter 20

Finite addresses, precision costs and physical interpretation

Finite combinatorics, metric error and physical coefficients can be assembled into a useful model, provided each has a declared role. An address may be perfectly legal while the represented piece remains large. A graph solve may be exact while its conductances fail to represent the intended medium. A plot may be visually convincing while its unresolved endpoint order changes the quotient. The aim of this chapter is constructive: state the interfaces between these tasks so that a valid estimate at one interface can actually be used at the next.

20.1 Legal addresses recognized by finite states

Fix a finite alphabet $\mathcal{A}$ of $b \geq 2$ symbols, a finite state set S of size s, an initial state s_0, and a partial deterministic transition function

$$\tau : S \times \mathcal{A} \longrightarrow S.$$

Partial means that some transitions are forbidden. Deterministic means that a state and a symbol have at most one next state. A finite word is legal when its successive transitions are all defined. Every state is accepted here; a specified set of final accepting states would replace the final vector of ones below by its indicator vector. Define the integer matrix

$$A_{ij} = \#\{a \in \mathcal{A} : \tau(i, a) = j\}.$$

An entry can exceed one because different symbols may have the same next state. The row sums are at most b.

Proposition 20.1 (Exact address counts)**.** *The number of legal words of length n starting at s_0 is*

$$N_n = e_{s_0}^{\mathsf{T}} A^n \mathbf{1} \leq b^n.$$

These counts can be obtained by integer additions along the transition table using $O(nsb)$ additions. Storing the counts after one level needs $O(s(n \log b + 1))$ bits, and the straightforward addition implementation uses at most $O(nsb(n \log b + 1))$ bit operations.

Proof. Let $v_k(j)$ count legal words of length k ending in state j. At level zero, $v_0 = e_{s_0}^{\mathsf{T}}$. A word ending at i extends to j in exactly A_{ij} ways, so $v_{k+1} = v_k A$. Induction gives $v_n = e_{s_0}^{\mathsf{T}} A^n$, and summing its entries gives the formula. Determinism ensures that distinct paths labeled by a given word cannot be counted twice. Without determinism the same matrix calculation counts paths, which need not count distinct words.

There are at most b^n length-n words in the full alphabet, proving the upper bound. To implement one step, initialize a zero count for each state and, for each defined transition (i, a), add $v_k(i)$ to its target count. There are at most sb such additions per step. Every intermediate count is at most b^{k+1} and therefore uses at most $1 + (k + 1)\log_2 b$ bits. Addition takes time linear in this number of digits by the usual carry algorithm. The stated total bound follows by bounding each of the n levels by its final digit length; keeping only the previous and next count vectors gives the storage bound.

The matrix formula can also be expanded as a finite sum over intermediate states, making its multiplicities explicit:

$$N_n = \sum_{i_1,\dots,i_n \in S} A_{s_0 i_1} A_{i_1 i_2} \cdots A_{i_{n-1} i_n} \qquad (n \geq 1).$$

Each factor counts allowed labels on one transition. Determinism means that once a word and starting state are fixed, the entire state sequence is forced; two terms cannot count that same word. For $n = 0$, the matrix formula separately gives $N_0 = 1$, the empty word. In an implementation, the counts for one step must all be computed from the old vector, not updated in place and reused as if they still belonged to the old level. Otherwise paths of different lengths can be mixed inadvertently. Summing the level-dependent digit costs yields the more explicit bound

$$O\left(sb \sum_{k=0}^{n-1} (1 + (k + 1)\log b)\right) = O\big(sb(n + n(n + 1)\log b)\big).$$

The two-vector storage bound follows from keeping at most $2s$ integers of final digit length. This is a count for exact integer additions and state-table access; constructing a justified automaton from a geometric system is not included in that arithmetic cost.

□

This calculation can be much smaller than listing every word. It does not construct the geometric pieces and does not test a claimed transition rule against the original puzzle. In a concrete application, the rule must be proved to include exactly the intended successors, or explicitly declared an over-approximation or under-approximation. An over-approximation can bound the number of possible words from above but may contain unrealizable itineraries. An under-approximation can omit the very boundary address one is trying to locate.

20.2 Depth, storage and geometric accuracy

A single length-n address in a b-symbol alphabet can be stored in $n\lceil \log_2 b \rceil$ bits by fixed-length symbol encoding, plus any needed record header. If one is given an exact child-membership oracle, following one address needs at most n successor choices. Such an oracle is a hypothesis: at a shared boundary, more than one child can contain the point, and from approximate coordinates the choice may be undecidable. An acceptable output can retain all compatible children instead of inventing a unique address.

With a mesh bound $\delta_n \leq Cq^n$, a localization tolerance $\varepsilon < C$ is met at

$$n_\varepsilon = \left\lceil \frac{\log(C/\varepsilon)}{\log(1/q)} \right\rceil. \tag{20.1}$$

Thus one address has storage proportional to $\log(1/\varepsilon)$ for fixed b, C, q. Complete enumeration can be much more expensive. For the full b-ary tree, its number of depth-n records is b^n, and the ceiling in (20.1) gives

$$(C/\varepsilon)^{\log b/\log(1/q)} \leq b^{n_\varepsilon} < b(C/\varepsilon)^{\log b/\log(1/q)}. \tag{20.2}$$

Writing one record per leaf already requires at least that many record writes. It is not a lower bound for every possible representation of the compact set: shared subtrees, exact symmetries, or a formula for the set can permit compression. It is a statement about the chosen complete leaf-list output.

Under the logarithmic modulus of Proposition 16.4, the required depth can instead grow like $\varepsilon^{-1/\beta}$, and full enumeration can grow exponentially in that quantity. A complexity claim about boundary localization must therefore name the modulus and the output representation. Saying only that "the pieces shrink" omits both.

Proposition 20.2 (Composing geometric and evaluation errors)**.** *Suppose $x \in K_v$, $\operatorname{diam} K_v \le \delta$, and $x_v \in K_v$ has an approximation $\widetilde{x}_v$ with $d(x_v, \widetilde{x}_v) \le \eta$. Let a scalar observable F be defined on both K and these approximate points, with modulus ν on that domain. If a calculated value y_v obeys $|y_v - F(\widetilde{x}_v)| \le \xi$, then*

$$|y_v - F(x)| \le \xi + \nu(\delta + \eta).$$

For a Lipschitz observable with constant M, this becomes $\xi + M\delta + M\eta$.

Proof. The triangle inequality gives $d(x, \widetilde{x}_v) \le \delta + \eta$. Apply the modulus to $F(x)$ and $F(\widetilde{x}_v)$, then add the calculation error ξ. The Lipschitz case substitutes $\nu(t) = Mt$. □

If F is known only on K, an approximate representative outside K cannot be substituted into it without an extension theorem or a different evaluation rule. This domain condition is often lost when error terms are assembled informally. Likewise, an asserted rounding bound for y_v must be justified for the actual evaluation procedure; it does not follow from the symbolic formula alone.

20.2.1 Allocating depth and coordinate precision separately

For a concrete error budget take $M = 5$, $\delta = 2^{-n}$, $\eta = 2^{-p}$ and $\xi = 10^{-4}$. The guaranteed error is

$$E(n, p) = 5(2^{-n} + 2^{-p}) + 10^{-4}.$$

For tolerance 10^{-2}, subtracting the fixed evaluation error and dividing by five gives the exact feasible condition

$$E(n, p) \le 10^{-2} \quad \Longleftrightarrow \quad 2^{-n} + 2^{-p} \le \frac{99}{50000}.$$

If n is fixed, define $T_n = 99/50000 - 2^{-n}$. No finite coordinate precision can work when $T_n \le 0$. For $T_n > 0$, the least nonnegative integer allowed by this bound is

$$p_{\min}(n) = \max\{0, \lceil \log_2(1/T_n) \rceil\}.$$

This solves a constrained error-allocation problem, rather than simply requiring the two contributions to be equal. All $n \le 8$ fail because already $2^{-8} > 99/50000$. At $n = 9$, $T_9 = 43/1600000$, which lies strictly between 2^{-16} and 2^{-15}; thus $p = 16$ passes and $p = 15$ fails. Balanced $n = p = 10$ passes since $2^{-9} < 99/50000$. The later worked choice $n = p = 11$ is a conservative valid choice with a larger margin. For equal unit costs per depth and precision bit, $n + p = 20$ is the least possible total: if $n + p \le 19$, then

$$2^{-n} + 2^{-p} \ge 2\sqrt{2^{-n-p}} \ge 2\sqrt{2^{-19}} > \frac{99}{50000},$$

where the final comparison can be checked by squaring positive numbers. The feasible choice $(10, 10)$ attains total twenty. Different computational costs for geometric refinement and coordinate precision can favor an asymmetric pair; the error inequality remains the constraint in either case.

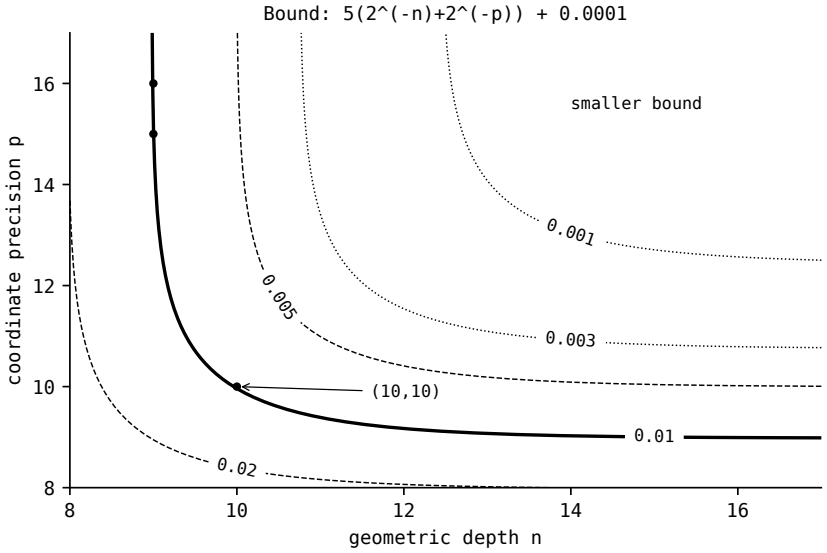


Figure 20.1: The explicit bound $E(n,p) = 5(2^{-n} + 2^{-p}) + 10^{-4}$ and its 10^{-2} threshold. Integer feasible pairs are selected from the displayed inequality; the drawn level sets are analytic-model illustrations, not error measurements from an executed geometric computation.

20.3 An exact one-dimensional conduction model

There is a direct physical setting in which the finite network equations are exact within the specified model. Represent every graph edge by an ideal one-dimensional rod of length $\ell_e > 0$, with scalar potential $U_e(x)$ and positive coefficient $k_e(x)$ on $[0, \ell_e]$. Assume k_e is piecewise continuous and bounded above and below by positive constants. For electric conduction one may take $k_e = \sigma_e A_e$, conductivity times cross-sectional area; for steady heat conduction use thermal conductivity times area. Impose the model law and steady conservation

$$J_e = -k_e(x) U_e'(x), \qquad J_e'(x) = 0$$

almost everywhere on each rod. Require U_e and J_e to be absolutely continuous, so the flux has no jump across a coefficient interface and steady conservation makes it constant on the whole rod. Also impose continuity of potential at ideal junctions, and conservation of the signed fluxes at every unforced junction. These are one-dimensional constitutive assumptions, not conclusions about an arbitrary planar hull.

Proposition 20.3 (Reduction of rods to a network)**.** *If the endpoint potentials of a rod are u at $x = 0$ and v at $x = \ell_e$, define*

$$R_e = \int_0^{\ell_e} \frac{\mathrm{d}x}{k_e(x)}, \qquad c_e = R_e^{-1}.$$

Then its flux in the increasing-x direction and its potential are

$$J_e = c_e(u - v), \qquad U_e(x) = u - J_e \int_0^x \frac{\mathrm{d}t}{k_e(t)}.$$

Its dissipation functional satisfies

$$\int_0^{\ell_e} k_e(x) |U_e'(x)|^2 \, \mathrm{d}x = c_e(u - v)^2.$$

Consequently the rod model at its junctions is exactly the finite Dirichlet network with conductances c_e.

Proof. Conservation makes J_e constant on each rod. Integrating $U_e' = -J_e/k_e$ gives the potential formula. Evaluating it at the far endpoint gives $u - v = J_e R_e$, and the positive finite integral R_e can be divided out. Substituting the derivative into the dissipation integral gives

$$J_e^2 \int_0^{\ell_e} \frac{\mathrm{d}x}{k_e(x)} = J_e^2 R_e = \frac{(u-v)^2}{R_e}.$$

At an unforced junction, summing outward fluxes yields $\sum_{y\sim x} c_{xy}(u_x - u_y) = 0$, exactly the graph harmonic equation. Summing the rod dissipation identities yields (18.1). Prescribed terminal potentials are the graph boundary values, so Theorem 18.1 gives their unique extension.

The same reduction has a weak variational derivation which handles coefficient jumps without differentiating k_e. On one rod, let U be the displayed solution and let W be absolutely continuous with $W(0) = W(\ell) = 0$ and finite energy. Since k is bounded below by a positive constant, $W' \in L^2$, and the constant-flux formula gives

$$\int_0^{\ell} kU'W' \,\mathrm{d}x = -J \int_0^{\ell} W' \,\mathrm{d}x = -J(W(\ell) - W(0)) = 0.$$

Expanding the energy of the competitor $U + W$ therefore yields

$$\int_0^{\ell} k|U' + W'|^2 \,\mathrm{d}x = \int_0^{\ell} k|U'|^2 \,\mathrm{d}x + \int_0^{\ell} k|W'|^2 \,\mathrm{d}x.$$

The last term vanishes only if $W' = 0$ almost everywhere, and absolute continuity with zero endpoints then gives $W = 0$. Thus the solution minimizes the continuous rod functional before any graph elimination. Conversely, suppose an absolutely continuous finite-energy potential satisfies the displayed weak equation for every such zero-endpoint test. Put $q = kU' \in L^2$ and $\bar{q} = \ell^{-1} \int_0^{\ell} q \,\mathrm{d}x$. The admissible test

$$W(x) = \int_0^x (q(t) - \bar{q}) \,\mathrm{d}t$$

has zero endpoints. Substitution shows $0 = \int q(q - \bar{q}) = \int (q - \bar{q})^2$, so q is constant almost everywhere. Taking $J = -\bar{q}$ recovers the flux formula and then the explicit potential by integration. This proves equivalence within the stated one-dimensional finite-energy class; it is not a derivation of that class from a multidimensional material model.

□

For constant coefficients, $c_e = k_e/\ell_e$. In the electric interpretation, σ_e has units siemens per metre, A_e has units square metres and ℓ_e has units metres, so c_e has units siemens. With potentials in volts, $c_e(u - v)$ is in amperes and $c_e(u - v)^2$ is in watts. Thus $\mathcal{E}_c$ is dissipated power in that interpretation, not electrostatic stored energy. For thermal conduction, c_e has units watts per kelvin, flux has units watts, and the quadratic functional has units watt-kelvin. Its minimization remains mathematically valid, but calling that quadratic quantity a stored heat energy would be dimensionally incorrect.

At an interface between two constant coefficients k_- and k_+, continuity of flux specifies the slope jump exactly:

$$k_- U_-' = k_+ U_+' = -J, \qquad \frac{U_+'}{U_-'} = \frac{k_-}{k_+} \quad \text{when } J \neq 0.$$

Potential stays continuous; its derivative need not. For $J = 0$ both slopes are zero and the ratio is not used. On finitely many constant segments with lengths ℓ_j and coefficients k_j, direct

integration gives $R = \sum_j \ell_j/k_j$. This is a weighted harmonic, not arithmetic, combination of coefficients. For arbitrary bounded positive k, another Cauchy–Schwarz calculation gives

$$\ell^2 = \left(\int_0^\ell \sqrt{k}\,\frac{1}{\sqrt{k}}\,\mathrm{d}x\right)^2 \le \left(\int_0^\ell k\,\mathrm{d}x\right) R, \qquad c \le \frac{1}{\ell^2}\int_0^\ell k\,\mathrm{d}x.$$

Equality requires k to be constant almost everywhere. The simpler pointwise bounds also give $k_{\min}/\ell \le c \le k_{\max}/\ell$. These inequalities explain why replacing a heterogeneous rod by its arithmetic mean coefficient can overestimate its conductance, even when all the lengths and material values are known exactly.

The rod model assumes no lateral leakage and ideal junctions without additional contact resistance. A real thin two- or three-dimensional domain requires a separate reduction theorem with geometric hypotheses, boundary conditions and an error bound. No such thin-domain limit is established here. In particular, drawing graph edges along legal paths in a continuum does not justify setting $c_e = 1$, nor does it prove that the network measures the continuum's conformal modulus or capacity.

20.3.1 Coefficient uncertainty before the graph solve

The reduction also identifies which physical uncertainties can be passed to the finite comparison theorem. Suppose the rod length is fixed and its true coefficient $\widetilde{k}$ obeys the pointwise relative bounds

$$(1-\eta)k(x) \le \widetilde{k}(x) \le (1+\eta)k(x), \qquad 0 \le \eta < 1.$$

Positive reciprocation followed by integration gives

$$\frac{R}{1+\eta} \le \widetilde{R} \le \frac{R}{1-\eta}, \qquad (1-\eta)c \le \widetilde{c} \le (1+\eta)c.$$

Thus a pointwise coefficient enclosure induces the same relative enclosure on the reduced conductance. When every rod has such an enclosure with the same η, Proposition 19.4 immediately bounds the two-terminal conductance and resistance of the entire network. This composition is deterministic; it makes no independence assumption and does not interpret the intervals as confidence levels.

For a constant electric coefficient $c = \sigma A/\ell$, separate positive parameter intervals can instead be propagated directly. If

$$\sigma \in [\sigma_-, \sigma_+], \quad A \in [A_-, A_+], \quad \ell \in [\ell_-, \ell_+], \qquad \sigma_-, A_-, \ell_- > 0,$$

then monotonicity of multiplication and division on positive numbers gives

$$\frac{\sigma_- A_-}{\ell_+} \le c \le \frac{\sigma_+ A_+}{\ell_-}.$$

This enclosure remains valid if the three parameters are correlated, although correlations may make it wider than necessary. A nominal value $c_0 > 0$ converts it to relative form by taking

$$\eta = \max\left\{1 - \frac{\sigma_- A_-}{c_0 \ell_+}, \frac{\sigma_+ A_+}{c_0 \ell_-} - 1, 0\right\},$$

provided the resulting η is less than one. If it is not, use the positive absolute conductance bounds directly in Rayleigh comparison; the failure of this convenient relative parameterization does not make the absolute intervals invalid.

These formulas distinguish three sources of error. Uncertain physical coefficients enter before reduction and alter the graph itself. A residual bound concerns the accuracy with which

the specified graph equations were solved. A discrepancy between the rod model and a real multidimensional sample is a model error requiring additional evidence. Adding more arithmetic precision can reduce the second contribution, but cannot eliminate either of the other two. Conversely, exactly known material coefficients do not excuse an unchecked linear-system solve.

There is an explicit finite-network estimate for the change of potential caused by coefficient uncertainty. For this estimate assume $I \neq \varnothing$; if there are no interior vertices, the potential-difference assertion is vacuous and no path constant R_* is needed. Let h solve the nominal network and $\widetilde{h}$ solve the perturbed one with the same boundary vector g. For $x \in I$, the residual obtained by inserting h in the perturbed equations is

$$s_x = (\widetilde{L}h)_x = \sum_{y \sim x} (\widetilde{c}_{xy} - c_{xy})(h_x - h_y).$$

Write $G = \max_B g - \min_B g$. The maximum principle and the relative coefficient bounds imply

$$|s_x| \leq \eta G \sum_{y \sim x} c_{xy}.$$

Use the same simple paths as for the nominal residual estimate. Their perturbed resistance sums are at most $R_*/(1-\eta)$, since $1/\widetilde{c}_e \leq 1/((1-\eta)c_e)$. Applying the already proved residual theorem to the perturbed network, with trial potential h, gives

$$\|h_I - \widetilde{h}_I\|_2 \leq \frac{mR_*}{1-\eta} \|s\|_2.$$

All constants refer to specified finite data, and the bound remains valid when coefficient errors are correlated. It can be conservative, especially when the nominal voltage drop on many edges is small. The exact formula for s_x retains those smaller drops, whereas the range bound G discards them for simplicity. Neither version includes a discrepancy between the one-dimensional constitutive model and a physical sample; that is still a separate modeling question.

20.4 What transfers from a finite model to a continuum

A precise transfer statement can nevertheless be useful. Let X be a compact space with a physical or analytic quantity $F : X \to \mathbb{R}$, let S_n be finite model states, and let $\iota_n : S_n \to X$ be a specified realization map. Suppose every $x \in X$ lies within distance δ_n of some $\iota_n(s)$, that F has a known modulus ν, and that model outputs $f_n(s)$ satisfy the independently justified comparison

$$|f_n(s) - F(\iota_n(s))| \leq \varepsilon_n \quad \text{for every } s \in S_n.$$

Then each selected nearby state gives

$$|f_n(s) - F(x)| \leq \varepsilon_n + \nu(\delta_n). \tag{20.3}$$

The proof is the triangle inequality, but the hypotheses identify the substantial work: a realization map, coverage, a modulus, and a comparison between the model and the intended quantity. A successful exact graph solve establishes none of these hypotheses automatically.

The topological and analytic sides can vary independently. Multiplying every conductance in a fixed graph by $\lambda > 0$ leaves its vertices, edges, cyclic order and all address labels unchanged. Its harmonic potentials for fixed boundary data remain unchanged because L is simply multiplied by λ, whereas

$$C_{ab}(\lambda c) = \lambda C_{ab}(c), \qquad R_{ab}(\lambda c) = \lambda^{-1} R_{ab}(c).$$

Thus even a perfect topological identification with a finite puzzle tree cannot determine a resistance without the coefficient scale.

Similarly, an exact finite noncrossing partition establishes only the order consistency of its supplied angle classes. A continuum pinched-disk model additionally involves closedness of the relation, the behavior of limit classes, and the stated hypotheses of the quotient theorem. Finite checking is valuable because it can reject a wrong finite proposal with a small exact witness. Its scope should remain visible when the proposal is accepted.

20.5 Solved investigations

Investigation 20.4 (Binary addresses without adjacent ones)**.** Count binary words of length n containing no consecutive ones using a two-state transition system. Give the count at $n = 5$.

Solution. Use state 0 when the previous symbol is zero or the word is empty, and state 1 when the previous symbol is one. From state 0, symbols $0, 1$ lead to states $0, 1$ respectively; from state 1, only symbol 0 is allowed and leads to state 0. Thus

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}, \qquad N_n = (1, 0)A^n \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

Splitting words by whether their first symbol is 0 or 10 gives $N_n = N_{n-1} + N_{n-2}$ for $n \geq 2$, with $N_0 = 1$, $N_1 = 2$. The resulting values are $1, 2, 3, 5, 8, 13$, so $N_5 = 13$. These are legal symbolic words; their geometric realizability still depends on the application supplying this transition rule.

Investigation 20.5 (One address versus a complete list)**.** A full binary model has $\delta_n \leq 2^{-n}$. Compare the storage of a single address with the number of leaf records required for tolerance 2^{-20}.

Solution. Depth 20 suffices. A single binary address takes 20 symbol bits, apart from its metadata. The complete level has $2^{20} = 1048576$ leaf records. If each record stores its full address separately, the symbol fields alone occupy $20 \cdot 2^{20}$ bits. A prefix tree can share those fields, and an implicit full-tree formula can avoid listing them at all; the million-record lower bound applies specifically to an explicit one-record-per-leaf output, not to every representation of the model.

Investigation 20.6 (An observable error budget)**.** Suppose F is Lipschitz with constant 5 on the relevant domain, the geometric mesh is at most 2^{-n}, representatives have error at most 2^{-p}, and evaluation error is at most 10^{-4}. Give integer choices ensuring total error at most 10^{-2}.

Solution. Take $n = p = 11$. Proposition 20.2 gives

$$|y_v - F(x)| \leq 10^{-4} + 10 \cdot 2^{-11} = 10^{-4} + 5/1024 < 10^{-2}.$$

This leaves a strict margin. The domain of the Lipschitz estimate must include the approximate representatives; otherwise the displayed application of the proposition is not licensed. A smaller actual mesh could allow fewer coordinate bits, but that tradeoff must still obey the same combined inequality.

Investigation 20.7 (Two materials in one rod)**.** A rod has length 3, with $k = 2$ on its first unit of length and $k = 4$ on its remaining two units. The endpoint potentials are 3 and 1. Compute its effective conductance, flux and interface potential.

Solution. The integrated resistance is $R = 1/2 + 2/4 = 1$, so $c = 1$ and the increasing coordinate flux is $J = c(3-1) = 2$. At the material interface, $U(1) = 3 - 2(1/2) = 2$. On the second segment the drop is another $2(2/4) = 1$, reaching the prescribed terminal value 1. The dissipation is $J^2R = 4$, agreeing with $c(3-1)^2$. Using an unweighted arithmetic mean of the coefficients in place of the integral of their reciprocals would not recover this result.

Investigation 20.8 (Why a correct graph solve is not a continuum proof)**.** A finite graph has been checked to have a legal noncrossing embedding, and its rational Dirichlet problem has been solved exactly. Does this prove that its effective resistance equals a chosen continuum capacity? Give a concrete obstruction and name the missing comparison data.

Solution. It does not. Multiply every graph conductance by 2. The embedding, noncrossing verdict and boundary potentials stay the same, while the resistance is divided by 2. Both graph solves are mathematically correct, so the topological information cannot determine which, if either, equals the proposed continuum quantity. One needs a specified normalization, a map or reduction connecting graph states to the continuum, and a proved comparison with an error bound. In the form (20.3), that missing comparison is the uniform bound $|f_n(s) - F(\iota_n(s))| \le \varepsilon_n$, together with coverage and a modulus. It is a mathematical obligation, not an empirical property established by a visually plausible graph.

The five chapters have supplied finite statements with complete written proofs, explicit hypotheses and exact worked examples. Their algebra does not certify all puzzle pieces of an unspecified continuum, an infinite lamination, or a physical experiment. Those broader claims would require the additional geometric and analytic interfaces just identified. The present statements are also not proof-assistant kernel checks; they provide a precisely scoped written foundation for such work.

Chapter 21

Source guide and precise obligations

This guide follows the hierarchy of Section 9 of Lyubich's supplied manuscript [1], printed pages 152–163. It records all 30 numbered units and their printed subparts. Definitions and substantive unnumbered arguments are included in their original subsection positions. The chapters preceding this guide instead follow proof dependencies. A cross-reference below identifies the precise statement actually established here; it does not automatically certify every stronger reading of its source counterpart.

The evidence level is a written mathematical argument with independent agent cross-review. Classical inputs are declared in the preface and dependency register. No entry below asserts Lean, Rocq, CAS or external peer-review acceptance. In particular, a correction, counterexample or remaining obligation is not counted as a proof of the literal source assertion.

21.1 9.1. Cut-curves and puzzle pieces for hulls

21.1.1 9.1.1. Terminology

The source distinguishes touching from cutting, allows finite intersections for general cut-curves, and distinguishes sectors, their shadows and perfect sectors. It also describes Green representatives with vertical ray portions and horizontal equipotential portions. Chapter 5 fixes the simple-cut convention and keeps a plane cell distinct from its intersection with the hull. The Green-coordinate formulas are developed in Chapter 2; access-preserving replacement is used explicitly in Proposition 6.3. A simultaneous ambient isotopy of an arbitrary infinite curve family is not claimed.

Rooted and unrooted branches differ by whether the root is included. Puzzle vertices and local corners do not by themselves establish connectedness of an entire hull intersection. Definition 5.5 therefore states the global connectedness and openness requirements used later. This is an explicit normalization, not a deduction from a drawing.

21.1.2 9.1.2. Chopping off subhulls

Lemma 9.1. Lemma 5.2 proves the full simple-cut statement: the two closed-side intersections are nondegenerate full continua, their union is the original hull, and their intersection is the root. The proof separates the connectedness argument from the use of fullness.

Corollary 9.2, parts (i) and (ii). Part (i) is proved for finite disjoint simple cut-lines in Proposition 5.3. Part (ii), if interpreted for every Jordan cut-curve allowed by the preceding broad definition, is false. Example 5.4 uses a U-shaped arc and a round disk whose intersection

is two disjoint segments. Admissible connected/full pieces supply the corrected setting for later arguments. Their admissibility must be proved in an application.

Remark 9.3. The finite-cell proof also identifies the possible common roots of cut-off subhulls. The claim concerns intersections with K. Two closures of entire plane sectors can share a separating curve, and should not be substituted for these subhulls.

Corollary 9.4, parts (i)–(iii). Theorem 6.1 proves (i) and the visible-root nondegeneracy condition (ii) for the stated disjoint simple cuts, with the base point avoiding their roots. Proposition 6.2 proves the nested compact/full intersection statement needed for (iii). Connectedness of the finite retained sets is a hypothesis or a prior conclusion. The unrestricted Jordan-piece assertion cannot be recovered by taking a limit after Corollary 9.2(ii) has already failed at a finite stage.

Lemma 9.5. Proposition 6.3 proves persistence and injectivity of accesses under visible-root removal. The retained roots witness both sides of a ray pair. The side-invariance argument preserves the retained subsets of K; inclusion then induces the stated injection of local approach classes. It need not be surjective: cutting can create new approaches.

The source's unnumbered sector-complement description is represented by the retained-hull intersection in Theorem 6.1. Its compactness, connectedness, fullness and possible degeneration are separate conclusions.

21.1.3 9.1.3. Puzzle ends, impressions, and rigidity

The source introduces puzzle fjords, escaping nests and impressions, and proposes an end attached to every point. Nested admissible sets do yield compact connected impressions. A universal assignment with the proposed fiber semantics is a stronger claim.

Exercise 9.6. The general equivalence between the source's escaping-end construction and pairwise nonseparation is not established here. The tangent-root example in Example 6.4 obstructs the proposed universal end assignment with disjoint fibers. Valid nested-intersection statements are proved in Proposition 6.2; they do not create a missing quotient or resolve every end convention.

Corollary 9.7, parts (i)–(iii). Part (i) holds for the admissible nests just specified. Literal part (ii) fails under open-side nonseparation: for two tangent disks the tangent point is inseparable from either center, while the centers are separated from one another. The relation is not transitive. Proposition 6.5 states sufficient genuine partition/closed-relation hypotheses. Under those hypotheses the singleton-fiber conclusion in (iii) follows by the compact to Hausdorff homeomorphism theorem. That conditional result does not repair the nonexistent original projection by assumption.

Proposition 9.8, parts (i)–(iii). Definition 7.1 distinguishes admissible and perfect rigidity. Proposition 7.2 proves weak local connectivity at an admissibly rigid point, local connectivity at a perfectly rigid point, and the corresponding global conclusion. Connected neighborhoods and connected open neighborhoods have different roles at a single point. Proposition 3.1 explains why the distinction disappears under the appropriate hypothesis at every point.

Corollary 9.9. The shrinking-diameter criterion follows from the same proposition and compact nesting. Proposition 7.3 explains why a filled interior body obstructs such shrinking for hull cuts. A singleton intersection supplies no predetermined geometric rate; Chapters 16 and 20 state what quantitative inputs do supply a rate.

21.1.4 9.1.4. Puzzle and rays

Exercise 9.10, parts (i)–(iv). Part (i) also needs an admissibility restriction. In the U-arc example, let $P = \overline{D}(i, 3/4)$ and take the horizontal crosscut $\sigma = [-1/2, 1/2] + i$. Its interior lies in the exterior of K and it lies inside P. The bounded fjord is the rectangle $(-1/2, 1/2) \times (0, 1)$, which contains $i/10 \notin P$. Thus crosscut containment in an arbitrary source piece does not imply fjord containment there.

For the impression-containment and subordinate-end assertions (ii) and (iii), Lemma 7.5 gives an actual finite-barrier statement away from its finitely many endpoint angles. Theorem 7.6 treats those exceptions in the shrinking situation. Merely touching a closed piece does not give arbitrary whole-impression containment. The general projection clause (iv) retains the unresolved end/quotient issue from Corollary 9.7; it is not claimed from the limited barrier theorem.

Corollary 9.11. Theorem 7.6 proves singleton impression and ray landing at an actual cut-rigid point using shrinking finite pieces and endpoint-angle control. Corollary 7.7 gives the complete landing-fiber description. This route proves the local conclusion without assuming a global puzzle quotient.

Remark 9.12. The source permits restricted puzzle families with an intersection-refinement property. That property is useful for choosing smaller pieces, but it is not asserted here to prove disjoint impressions or a closed quotient relation. The separate requirements of Proposition 6.5 remain in force.

21.1.5 9.1.5. Branches and limbs

The source introduces well branching, pointed hulls, wakes, limbs and bodies. The text keeps a full hull, a chosen body, a sector and its attached continuum distinct. Chosen internal centers are additional data.

Corollary 9.13. Proposition 10.4 proves connected consecutive sectors for a finite landing fiber in a locally connected hull, with at least two accesses. The condition $m \geq 2$ matters: an endpoint of an arc can have a single access and does not meet a definition requiring branching valence at least two. The extension to a merely rigid point in a globally non-locally-connected hull remains a separate obligation.

Lemma 9.14. The source's full criterion using approximating pairs of landing angles to prove connectivity of a sector branch remains unproved under its weakest stated hypotheses in this edition. Neither the null-family lemma nor the continuous global landing-map argument is substituted for that theorem. This entry preserves the exact missing source unit for subsequent work.

Exercise 9.15. Proposition 7.8 gives a finite-cover local-connectivity criterion when the selected open cut cells are connected. To recover the source's full perfect-cut-rigidity conclusion one must additionally construct compatible small Jordan pieces and verify their corners. That construction is not inferred from pairwise separation alone and remains an obligation.

21.2 9.2. Interior components of hulls

21.2.1 9.2.1. Cut-curves and puzzle pieces for general continua

The source enlarges cuts to internal Green arcs and dipoles meeting a continuum at two points. The number of contact points changes the topology; the simple-cut proof cannot be copied without checking it.

Exercise 9.16. The singleton-cut connectedness proof extends to arbitrary continua without the fullness conclusion. The general dipole claim is false: Example 9.1 intersects the unit circle with the band between two disjoint horizontal dipoles, leaving two separate arcs. Additional admissibility is therefore necessary for multiple dipole cells.

Exercise 9.17. Definition 5.5 requires connectedness of the relevant open intersection explicitly for a perfect piece. A proof from an ambiguous local-corner condition alone is not claimed. Recovering that implication under a precisely specified source corner definition remains an obligation.

Corollary 9.18. Under that explicit perfect-piece definition, the local-connectivity proof of Proposition 7.2 applies to an arbitrary continuum J. It uses connected relatively open neighborhoods, not fullness of J.

21.2.2 9.2.2. Limbs and local connectivity at peripheral points

The unnumbered Jordan-body decomposition in source equation (9.1) motivates Definition 9.2. It specifies the attachment points, pairwise intersection rules and null-family condition.

Lemma 9.19, parts (i)–(iii). Proposition 9.4 proves local connectivity at nonroot body-boundary points for a null family of attached continua. Proposition 9.5 proves the stronger global gluing statement under local connectivity of the limbs, with its stated access consequence. These are scoped results related to (i). The full source access assertion under weaker hypotheses is retained as an obligation.

The perfect rigidity in (ii) must distinguish $J = \partial K$ from the filled hull K: the proposed small dipole puzzles concern J, whereas an interior body prevents hull-cut shrinking. Part (iii), the full inherited root-rigidity construction, likewise requires a boundary puzzle argument not supplied by the gluing theorem. The missing peripheral cut constructions are recorded, not treated as consequences of a continuous boundary map by terminology.

Lemma 9.20. Proposition 9.6 proves that dense roots force shrinking in the explicit Jordan-body decomposition with disjoint Green-root sectors. Its two-root dipole barrier uses the body and sector hypotheses. It is not the false general inference that bounded area forces every disjoint continuum to have small diameter.

21.2.3 9.2.3. LC hulls and Jordan disks

Proposition 9.21. Theorems 8.1 and 8.3 supply the gate retraction and the Jordan-domain proof for every interior component of a locally connected full continuum. The boundary is shown locally connected through a continuous image of the exterior boundary, rather than by assuming that an arbitrary closed subset of a locally connected space is locally connected. Injectivity of the interior boundary map is then proved using fullness.

Exercise 9.22. Proposition 8.2 proves that two distinct interior component closures meet in at most one point. Continuity and constancy of the gate on an outside component provide the argument.

21.3 9.3. Legal issues

The source defines pointed components, legal points and arcs, legal convexity, and the finite legal hull. Definition 10.1 constrains the intersection with each *closed* interior component. An open-component-only condition permits arbitrary boundary semicircles in a disk, so it cannot give uniqueness. The closed-component convention is an explicit normalization of the objects studied here.

Exercise 9.23. Theorem 10.2 proves existence and uniqueness of a regulated legal joining arc. The existence proof controls infinitely many replacements using the null family of component diameters; parameter-interval lengths alone would not ensure continuity.

Exercise 9.24. Proposition 10.3 proves the separating-access statement for an interior boundary point on a regulated legal arc. The external ray pair separates the two arc branches. The endpoint and interior-point conditions remain distinct.

Exercise 9.25. Proposition 10.5 proves the origin-crossing conclusion when the pointing is equivariant under $z \mapsto -z$. Example 10.6 uses a unit disk pointed at a nonzero imaginary internal center to show why geometric symmetry of K alone is insufficient. If the source's symmetry convention includes the pointing, this is a hypothesis clarification rather than a refutation of that reading.

Exercise 9.26. Theorem 10.7 proves that the finite regulated legal hull is a topological tree, gives its canonical joining paths and establishes the leaf/branch bound. Its topology supplies no electrical weights automatically.

21.4 9.4. Pinched disk model for a hull

21.4.1 9.4.1. Locally connected case

The unnumbered construction starts from the external landing equivalence and extends it over convex-hull classes. Chapters 11–13 prove the closed-relation and hull-extension steps, including singleton, two-point and universal classes.

Lemma 9.27. Theorem 12.6 proves that distinct landing fibers are unlinked; Lemma 12.5 connects circular order to disjoint hyperbolic convex hulls. The proof is distinct from an inspection of finitely many plotted chords.

Theorem 9.28, main statement and plane extension. Theorems 13.3 and 14.3 construct and recognize the pinched model, including its extension over the plane. Moore's sphere theorem is a named classical input with its hypotheses checked. The matching of interior components and continuity across their accumulation are separate steps. The latter uses Proposition 3.5, proved from the gate modulus, followed by the pasting argument. This repairs the source's mismatched reference to "Proposition 9.27" in its shrinking-component step; the preceding source item is an unlinking lemma, not that null-family result.

21.4.2 9.4.2. General case

The source proposes closing and cleaning cutpoint geodesics and distinguishing black and white gaps. Those words alone do not specify a unique closed equivalence relation, its decomposition classes or the required projection. Chapter 15 shows why limits can lose transitivity and separates minimal closed completion from planar nonseparation and local connectivity.

Proposition 9.29, projection and recognition clauses. Theorem 15.3 proves a well-posed conditional statement for an explicitly supplied closed monotone relation with the specified quotient properties. Proposition 15.2 constructs the least closed equivalence completion, but does not assert that this completion automatically has nonseparating classes or a locally connected quotient. Defining the source's general cleanup, constructing its claimed projection and proving its recognition criterion remain substantive obligations. Density of selected accesses does not discharge them.

Exercise 9.30. Proposition 15.5 and Theorem 15.6 prove that every continuous monotone map from the specified Cantor fan onto a locally connected compact Hausdorff continuum is constant. Thus its universal monotone locally connected model is a point in that precise category. Identifying this model with the still-unspecified construction of Proposition 9.29 remains unproved. The finite-fan comparison and the radius-map counterexample explain both perfectness of the angle set and the connected- fiber hypothesis.

What is ready for subsequent formalization

The canonical statement register records the exact authored statements and their surrounding hypotheses. Compactness lemmas, the finite circular-order tests, tree quadratic identities and residual bounds provide bounded starting points. The topology inputs and every qualification above remain visible in that programme. A successful build, a correct figure or a checked finite example is not a substitute for any of the remaining general proof obligations.

Bibliography

[1] Mikhail Lyubich, *Conformal Geometry and Dynamics of Quadratic Polynomials*, Volumes I–II, supplied 702-page manuscript, PDF metadata 14 September 2023. Principal source: §9, pp. 152–163, ending before §10 on p. 163. `https://www.math.stonybrook.edu/~mlyubich/book.pdf`. Page references concern the supplied file identified in the source record; no byte identity with a later online edition is asserted.

[2] Vladlen Timorin, *Moore's theorem*, arXiv:1001.5140v1 (2010). `https://arxiv.org/abs/1001.5140`. Author-hosted manuscript. Theorem 1.1 supplies the sphere-decomposition input; Lemma 2.5 supplies the path-to-arc input. These named classical results are not reproved here.

[3] Christopher J. Bishop, *Boundary Interpolation Sets for Conformal Maps*, author manuscript dated 25 August 2004. `https://www.math.stonybrook.edu/~bishop/papers/moore.pdf`. Section 3, Theorem 9 and Lemmas 12–13 provide context for closed decompositions and convex-hull extensions. The compactness proof used in this volume is given in full, including degenerate classes.

[4] Peter G. Doyle and J. Laurie Snell, *Random walks and electric networks*, version dated 5 July 2006, derived from the 1984 MAA Carus Monograph. `https://math.dartmouth.edu/~doyle/docs/walks/walks.pdf`. Sections 1.3.5–1.4.1 are the inspected source for classical network principles; this volume supplies finite quadratic-form derivations.

[5] Danny Calegari, *Topics in Point Set Topology*, author notes dated 7 March 2025. `https://math.uchicago.edu/~dannyc/notes/point_set_topology.pdf`. Supplementary context for the local-connectivity and path-construction architecture. This is not a claim to have checked every result in the notes.

[6] PHI research collective, *Hyperbolic Geometry and Holomorphic Contraction*, PHI Technical Monographs, Volume VII. Series background on the hyperbolic metric and planar Riemann mapping.

[7] PHI research collective, *Prime Ends and Conformal Boundary Geometry*, PHI Technical Monographs, Volume VIII. Series background on prime-end impressions and Carathéodory extension.