

The Continuation Geometry of Circuits

Flyxion
Independent Researcher

2026

Abstract

This book develops a falsifiable geometric and continuation-theoretic reading of circuit theory, in the same disciplined register as this author's prior treatment of locomotion. Impedance is shown to be, exactly rather than by analogy, an amplitwist operator in Needham's sense; two-port cascades compose amplitwist operators under a matched-impedance condition stated precisely; the diode is shown to require non-invertibility as a matter of necessity, not convenience, for any device performing rectification; capacitors, inductors, digital memory, and phase-locked loops are shown to be defined by accumulated history rather than instantaneous state, with a proof rather than an assertion at the center of that claim; injection locking and the Kuramoto model are shown to be independently convergent instances of the same equation, decades apart, in two unrelated engineering and biological literatures; and Johnson–Nyquist noise is shown to tie dissipation and fluctuation together by an independently established physical theorem. Explicit admissibility conditions, stated failure modes, and a closing chapter of open problems hold the whole framework to the standard that it must be capable of being wrong. The book closes by declining to assume its own further generalization beyond circuits and locomotion, on the same grounds it uses throughout: a unifying claim is worth making only if it specifies, in advance, what would count against it.

Contents

1	Charge, Field, and the Scalar Potential Φ	6
2	Kirchhoff's Laws as Conservation Principles	6
3	Passive Elements as Constraint Operators	7
4	Admissibility of a Circuit State	7

5	Phasors as Amplitwist Operators	9
5.1	The Claim, Stated Without Interpretation	9
5.2	Operator Addition Is Not Operator Composition	10
5.3	Resonance as the Vanishing of Twist	11
5.4	What This Chapter Does Not Yet Give	11
5.5	Verdict	12
6	Two-Port Networks as Amplitwist Chains	13
6.1	Cascading Is Matrix Composition, Not Scalar Composition	13
6.2	Where the Scalar Picture Is Recovered Exactly	14
6.3	Assembly Index of a Two-Port Cascade	15
6.4	Where the Scalar Picture Breaks Down	15
6.5	Verdict	16
7	Resonance as an Amplitwist Fixed Point	17
7.1	Resonance in the Standard Vocabulary	17
7.2	The Twist of an Impedance Operator	17
7.3	Fixed Points of Operator Action	18
7.4	Resonance and Admissibility	19
7.5	Resonance and Synchronization	19
7.6	Where the Fixed-Point Picture Breaks Down	19
7.7	Verdict	20
8	Filters as Admissibility Gates in Frequency Space	20
8.1	A Filter as a Selection Operator	20
8.2	Admissibility in Frequency Space	21
8.3	Low-Pass Filters as Continuation Preservers	21
8.4	Band-Pass Filters and Admissibility Windows	21
8.5	Stopbands as Refusal Regions	22
8.6	Bandwidth as Admissibility Volume	22
8.7	Transition Regions and Boundary Effects	22
8.8	Resonance and Filtering	23
8.9	Where the Admissibility Picture Becomes Incomplete	23
8.10	Verdict	23
9	The Diode as an Irreversible Spherepop Operator	23
9.1	Recalling the Pop Operator	24
9.2	The Ideal Diode as Pop Paired with Refuse	24
9.3	Irreversibility Is Not an Incidental Feature: Rectification Requires It	25
9.4	Dissipation: Where Entropy Enters This Book	26
9.5	Where the Two-State Model Breaks Down	26
9.6	Verdict	27

Histories Before States	28
10 BJT/FET as Constrained Amplification Fields	32
10.1 The Active Region Is Invertible	32
10.2 Cutoff and Saturation: Where Collapse Returns	33
10.3 The FET: A Different Nonlinearity, the Same Structure	33
10.4 Verdict	34
11 Op-Amps as Idealized Admissibility Engines	34
11.1 The Golden Rules as Contextual Admissibility	34
11.2 The Integrator: Where the Histories Chapter Becomes a Design Choice	35
11.3 The Comparator: Where Admissibility Gives Way to Collapse	35
11.4 Verdict	35
12 Log/Exponential Amplifiers as Compression Operators	36
12.1 The Log Amplifier as Dynamic-Range Compression	36
12.2 The Log-Antilog Multiplier and Chapter 6's Composition Law	36
12.3 Where This Chapter's Idealizations Fail	37
12.4 Verdict	37
13 Feedback as Recursive Continuation	38
13.1 The Closed Loop as a Fixed-Point Equation	38
13.2 Feedback in Time: Self-Reference as an ODE	38
13.3 A Brief, Deferred Note on History	39
14 The Barkhausen Criterion as an Admissibility Condition	39
14.1 The Criterion, Recalled	39
14.2 The Criterion Is the Linear Onset Condition for a Hopf Bifurcation	39
14.3 Necessary, Not Sufficient: Startup Requires More	40
14.4 Admissibility, Named Precisely	40
14.5 Where This Chapter's Picture Breaks Down	41
14.6 Verdict	41
15 Oscillator Circuits as Self-Sustaining Amplitwist Loops	42
15.1 The General Structure, Recalled and Specialized	43
15.2 The RC Phase-Shift Oscillator	43
15.3 The LC (Hartley) Oscillator	44
15.4 The Crystal Oscillator	44
15.5 Amplitude Self-Limiting, Concretely	45
15.6 Where This Chapter's Idealizations Break Down	45
15.7 Verdict	45
Synchronization: From Coupled Oscillators to Network Coherence	46

16	The Phase-Locked Loop as Engineered Continuation	49
16.1	The Loop, in Block Form	50
16.2	The Simplest Case Is Exactly Adler’s Equation	50
16.3	What Engineering the Coupling Buys: An Integrator	51
16.4	A Second Extension: Capture Range	52
16.5	Failure Modes Specific to This Chapter	52
16.6	Verdict	52
17	Digital Logic as Spherepop Collapse	53
17.1	Every Standard Logic Gate Is a Pop Operator	53
17.2	Not All Collapses Are Equal	53
17.3	Landauer’s Bound, Properly Earned Here	54
17.4	Reversible Gates: The Exception That Proves the Rule	55
17.5	Sequential Logic: Where History Re-enters	55
17.6	Failure Modes Specific to This Chapter	55
17.7	Verdict	56
18	Clock Distribution, Skew, and Domain Crossing	56
18.1	Clock Skew Is Deterministic, Not Statistical	56
18.2	Setup and Hold as a Joint Admissibility Criterion	57
18.3	Clock Domain Crossing: Where Chapter 17’s Metastability Is Un- avoidable	58
18.4	Distributed Synchronization: A Direct Application of an Earlier Chap- ter	58
18.5	Failure Modes and Limitations	59
18.6	Verdict	59
19	Memory and Repair	59
19.1	DRAM as Continuously Repaired History	60
19.2	Error-Correcting Memory as the Repair Chain, Concretely	60
19.3	Silent Miscorrection: When Repair Produces a Wrong, Admissible History	61
19.4	Failure Modes and Limitations	62
19.5	Verdict	62
20	Noise as Entropic Production	62
20.1	Entropy Production in a Dissipative Circuit	63
20.2	The Same Resistors That Dissipate Also Fluctuate	63
20.3	Shot Noise: A Second, Distinct Mechanism	64
20.4	Reactive Elements Are, Ideally, Silent	64
20.5	Where This Chapter’s Picture Is Incomplete	65
20.6	Verdict	65

21	Signal-to-Noise Ratio and the Compactness Criterion	65
21.1	Signal and Noise as Competing Components	66
21.2	The Recoverability Interpretation	66
21.3	Distinguishability and Admissibility	66
21.4	Noise as Compression Pressure	67
21.5	Shannon Capacity as an Admissibility Volume	67
21.6	The Compactness Criterion	68
21.7	A Circuit-Theoretic Example	68
21.8	Relation to Earlier Chapters	68
21.9	Where the Compactness Picture Breaks Down	68
21.10	Verdict	69
22	Failure Modes and Limitations	69
22.1	What This Book Actually Claims	69
22.2	Failure Mode I: Idle Renaming	70
22.3	Failure Mode II: Breakdown of Scalar Amplitwist Structure	70
22.4	Failure Mode III: Non-Sinusoidal Regimes	70
22.5	Failure Mode IV: Adaptive Admissibility	71
22.6	Failure Mode V: Lock Loss and Synchronization Collapse	71
22.7	Failure Mode VI: Entropy Without Dissipation	71
22.8	Failure Mode VII: History Without Continuation	72
22.9	A Compactness Criterion for the Framework Itself	72
22.10	Summary	72
23	Two-Port Cascades and Assembly Index: A Comparative Note	73
23.1	Two Different Composition Operations, One Shared Name	73
23.2	A Natural, Unresolved Extension	74
23.3	Verdict	74
24	Open Problems	74
24.1	The Full Matrix Case	74
24.2	Nonequilibrium Entropy and the Diode	75
24.3	Non-Periodic and Transient Behavior	75
24.4	Reversible and Adiabatic Circuits	75
24.5	Correlated, Multi-Cell Failure and Repair	75
24.6	Admissibility as a Design Methodology	76
24.7	Beyond Circuits	76
	Appendix: A Classification of This Book's Claims	76

1 Charge, Field, and the Scalar Potential Φ

This book has, since Chapter 5, used Φ , v , and S as labels connecting circuit quantities to this author's broader work on scalar-vector-entropy fields. It has not yet stated, formally, what licenses the connection at the level of circuits specifically, rather than assuming it. This chapter supplies that grounding, and states plainly where the connection is exact and where it is only a shared vocabulary.

Definition 1 (Node potential as Φ). *The voltage at a circuit node is the electric scalar potential Φ at that point, defined, as in ordinary electromagnetism, by $\Phi(\mathbf{r}) = -\int \mathbf{E} \cdot d\mathbf{l}$ along any path to a chosen reference point.*

Observation 1. *This identification requires no argument, in the same sense Chapter 5's identification of impedance with an amplitwist operator required none: voltage already is a scalar potential, by the ordinary definition of the term in physics, independent of anything in this book. What this book adds is only the decision to track it, alongside current and dissipation, using the same three-letter vocabulary (Φ, v, S) used elsewhere in this author's work on scalar-vector-entropy fields more generally. This is a shared labeling convention across domains, not a claim that a circuit is a literal instance of any cosmological field theory developed under the same name elsewhere; the correspondence asserted in this book is formal and structural (both are scalar potentials, vector flows, and entropy fields playing analogous roles), not an identity of physical domains.*

Observation 2. *Like any scalar potential, Φ is defined only up to an additive constant: choosing a reference ("ground") node fixes that constant by convention, not by any physical necessity. This is exactly the gauge freedom familiar from other scalar potentials in physics, and circuit ground should be understood as a choice of reference, not as an absolute zero the physical system itself distinguishes.*

The vector field v of this book's recurring (Φ, v, S) vocabulary is current, or more precisely current density: the physical flow whose divergence-free circulation (in the steady state) Section 2 formalizes as Kirchhoff's current law.

2 Kirchhoff's Laws as Conservation Principles

Observation 3. *Kirchhoff's current law (KCL), that the currents into a node sum to zero, is not an independent postulate of circuit theory; it is charge conservation, $\nabla \cdot \mathbf{J} = 0$ in the steady state, applied at a point. No charge accumulates at an ordinary circuit node, so whatever flows in must flow out. Kirchhoff's voltage law (KVL), that voltage drops around a closed loop sum to zero, is likewise not an independent postulate; it is the statement that Φ is a genuine scalar potential, since a conservative field has zero circulation around any closed path, $\oint \mathbf{E} \cdot d\mathbf{l} = 0$.*

Observation 4. *KVL's derivation from Φ 's conservative character depends on a quasi-static assumption: that any time-varying magnetic flux linking the loop is negligible. Where this assumption fails, most relevantly at high frequencies or in circuits with significant loop-enclosed changing flux, KVL as ordinarily stated no longer holds exactly, and the fuller Maxwell–Faraday treatment is required. This is not a defect in the reading offered here; it is the same kind of boundary this book has named for every other idealization it has relied on, from ideal losslessness in Chapters 5–6 to the equilibrium assumption behind Chapter 20's fluctuation-dissipation theorem.*

3 Passive Elements as Constraint Operators

Chapters 5, 6, and the Histories-before-States chapter, later in this book, each treat resistors, and then capacitors and inductors, in detail. This section states, once and in general, the organizing distinction those later chapters each specialize.

Definition 2 (Memoryless versus historical constraint). *A passive element imposes a memoryless (algebraic) constraint between its voltage and current if that constraint involves no integral or derivative with respect to time, and a historical (integral) constraint if it does.*

Observation 5. *A resistor's constraint, $v = iR$, is memoryless: the current at an instant is fixed by the voltage at that same instant alone. A capacitor's and an inductor's constraints, $v_C(t) = v_C(t_0) + \frac{1}{C} \int_{t_0}^t i d\tau$ and $i_L(t) = i_L(t_0) + \frac{1}{L} \int_{t_0}^t v d\tau$, are historical in exactly the sense the Histories-before-States chapter later proves rigorously for the capacitor case. This section's contribution is only to name the general taxonomy, memoryless versus historical constraint, before Chapters 5–6 and the Histories chapter specialize it; those later chapters should be read as worked instances of a distinction introduced here once, rather than as introducing it independently each time.*

4 Admissibility of a Circuit State

Four chapters written before this one, Chapter 8 (filters), Chapter 14 (the Barkhausen criterion), Chapter 19 (memory and repair), and Chapter 21 (signal-to-noise ratio), each used the word “admissible” with its own local, context-specific meaning, and none of them stated a general definition first. This section supplies the missing general definition and shows that all four prior uses are instances of it, not four independent borrowings of a convenient word.

Definition 3 (Physical admissibility). *A configuration of voltages and currents (or a trajectory of such configurations) is physically admissible if it satisfies Kirchhoff's current and voltage laws together with the constitutive relation of every branch element in the circuit, i.e. if it is a genuine solution of the circuit's governing equations.*

Definition 4 (Contextual admissibility). *Given a physically admissible configuration and a context-specific criterion C , relevant to whatever question is being asked of the circuit, the configuration is contextually admissible relative to C if it additionally satisfies C .*

Observation 6. *Every prior use of “admissible” in this book is contextual admissibility relative to a specific C , built silently on top of physical admissibility, which is assumed throughout and never separately argued for because a configuration violating Kirchhoff’s laws would not be a physically realizable circuit state to begin with. Chapter 8’s frequency admissibility takes C to be membership in a passband; Chapter 14’s oscillation admissibility takes C to be satisfaction of the Barkhausen magnitude and phase conditions; Chapter 19’s memory admissibility takes C to be passing a validation check (a parity syndrome, or a voltage remaining above threshold); Chapter 21’s detection admissibility takes C to be a signal-to-noise ratio exceeding a stated threshold. Four apparently unrelated uses of the same word turn out to share one structure, stated here for the first time rather than four times independently.*

Proposition 1 (Admissibility can be empty). *For some circuits and some contextual criteria C , no physically admissible configuration also satisfies C ; contextual admissibility relative to C is then the empty set.*

Proof. A purely passive network (composed only of positive resistances and ideal, lossless reactive elements, with no active gain element) cannot satisfy the Barkhausen magnitude condition of Chapter 14 with sustained, non-decaying amplitude, since a passive network’s loop gain magnitude cannot exceed unity without an external energy source, and equality alone (achieved only in the lossless limit) gives neutral stability rather than the self-starting growth Chapter 14, later in this book, shows sustained oscillation actually requires. No configuration of a purely passive network is therefore contextually admissible relative to “sustained self-starting oscillation,” regardless of how its components are chosen. \square

Observation 7. *This is an important, and previously unstated, limit on the whole admissibility vocabulary used throughout this book: asking whether a configuration is admissible presupposes that some physically admissible configuration could, in principle, satisfy the criterion in question. A criterion that is unsatisfiable in this way is not a hard case for the framework to work through; it is a sign that the criterion itself was badly posed relative to the physical resources actually available, a distinction worth keeping in mind before treating every empty admissibility set as a puzzle to be solved rather than a question to be reformulated.*

Verdict

These four chapters do not introduce new circuit phenomena; every specific fact used here (voltage as scalar potential, KCL and KVL as conservation laws, the memoryless/historical taxonomy, and the four prior chapters’ admissibility notions) was

already, in some form, present or presupposed elsewhere in this book. Their contribution is retroactive: they give this book’s recurring vocabulary, Φ , Kirchhoff’s laws, constraint taxonomy, and above all admissibility, a stated general form for the first time, and show that four chapters written without that general form in hand were nonetheless using a single, consistent structure rather than four unrelated senses of a convenient word. Proposition 1’s emptiness result is the one genuinely new technical content in this unit, and it is a boundary statement in the same register as every other failure mode named in this book: admissibility is a meaningful question only relative to criteria the underlying physics can actually satisfy.

5 Phasors as Amplitwist Operators

This chapter is written as a pilot, in the sense argued for before any other chapter of this book: if the claim that circuit impedance is an amplitwist operator cannot be made to do real conceptual work, the rest of the project needs revision before it needs more chapters. The chapter therefore holds itself to a stricter and more skeptical standard than a normal chapter would, and closes with an explicit verdict rather than a summary.

5.1 The Claim, Stated Without Interpretation

Under sinusoidal steady-state analysis, a linear circuit element relates the phasor voltage across it to the phasor current through it by

$$V(j\omega) = Z(j\omega) I(j\omega), \quad Z(j\omega) = |Z(j\omega)| e^{j\theta(j\omega)},$$

where $|Z|$ is the magnitude and θ the phase of the impedance. This is not a reinterpretation of anything. It is the ordinary definition of impedance, taught in every introductory circuits course, and the polar decomposition above is standard practice, not a step this chapter is introducing.

Definition 5 (Impedance as amplitwist operator). *Under the identification of Section 3.1 of this author’s prior work (Needham’s decomposition of a local holomorphic action into a scaling r and a rotation θ), the impedance $Z(j\omega) = |Z(j\omega)| e^{j\theta(j\omega)}$ acting on the current phasor $I(j\omega)$ to produce the voltage phasor $V(j\omega)$ is an amplitwist operator with amplitude $|Z(j\omega)|$ and twist $\theta(j\omega)$.*

Observation 8. *Unlike the identification made for the Kuramoto order parameter in this author’s work on locomotion, this identification requires no argument. A complex number acting by multiplication on another complex number is, definitionally, a scaling and a rotation; calling Z an amplitwist operator is true by the ordinary meaning of complex multiplication, not by an inference that needed to be established. The question this chapter must answer is therefore not whether the identification holds — it holds trivially — but whether it is worth making at all.*

This observation is the reason the chapter is a pilot rather than a formality. A relabeling that is true but idle would not justify a book; it would justify a footnote. The remainder of the chapter looks for a place where the relabeling changes what a reader notices, not merely what a reader calls it.

5.2 Operator Addition Is Not Operator Composition

The first candidate for genuine work done by the amplitwist framing is a distinction that ordinary circuit pedagogy leaves implicit and that the operator vocabulary makes forced rather than optional.

Proposition 2 (Series and parallel combination are additive, not compositional). *Let Z_1 and Z_2 be amplitwist operators (impedances) carrying the same current I in series. Then*

$$V_{total} = Z_1 I + Z_2 I = (Z_1 + Z_2) I,$$

so the series combination is the operator $Z_1 + Z_2$, obtained by adding the two operators, not by applying one operator's output as the other's input. Dually, for two admittances Y_1, Y_2 in parallel across the same voltage V ,

$$I_{total} = Y_1 V + Y_2 V = (Y_1 + Y_2) V,$$

again an addition of operators acting on a common input, not a composition of one operator's output feeding the next operator's input.

Proof. Immediate from Ohm's law and Kirchhoff's laws: in series, the same current is shared and voltages add; in parallel, the same voltage is shared and currents add. No step in either derivation applies Z_2 to the output of Z_1 , which is what operator composition, in the sense of Section 4 of this author's work on hierarchical operator chains, would require. \square

Observation 9. *This is the first place the amplitwist vocabulary earns something beyond restatement. Ordinary circuit pedagogy teaches series and parallel combination as algebraic rules to be memorized (impedances add in series, admittances add in parallel) without naming what kind of mathematical operation combination is, as opposed to what cascading two circuit stages is. Once impedance is named as an operator, the question becomes forced: is combining two impedances the same kind of operation as chaining two circuit stages together? Proposition 1 shows the answer is no. Series and parallel combination are additive operations on operators sharing a common input; cascading is compositional, in the sense of one operator's output becoming the next operator's input, and is deferred to two-port cascades in Chapter 6, where ABCD matrices multiply rather than add under cascading. Naming impedance an operator forces this distinction into the open; ordinary phasor pedagogy does not require it to ever be stated.*

This is a genuine, if modest, conceptual clarification: it is available in ordinary complex-number terms without the word “amplitwist,” but in practice it is rarely stated as a distinct fact in introductory treatment, where series and parallel rules are typically presented as two unrelated memorized formulas rather than as two instances of the same underlying operation (operator addition) applied to two different circuit topologies, contrasted explicitly with a third, different operation (operator composition) reserved for cascading. The operator vocabulary does not discover new mathematics here; it makes an existing but usually implicit structural distinction impossible to skip past.

5.3 Resonance as the Vanishing of Twist

The second candidate is a reframing of a standard result rather than a new one, offered with the same honesty about its modesty.

Proposition 3 (Resonance as pure scaling). *A series RLC circuit is at resonance when its reactive components cancel, $\omega L = \frac{1}{\omega C}$, at which frequency $Z(j\omega) = R$ is purely real. In amplitwist terms, resonance is the condition under which the twist $\theta(j\omega)$ vanishes: the operator degenerates from a genuine rotation-and-scaling to a pure scaling, and voltage and current phasors become collinear.*

Observation 10. *This is not a new result; “impedance is purely resistive at resonance” is the standard statement, and this chapter is not claiming otherwise. What the amplitwist framing offers is a different, and arguably more geometrically transparent, way of holding the same fact in mind: resonance is not a special value of a formula, but the specific condition under which an operator that is generically a rotation-and-scaling becomes a scaling alone. Whether this reframing is pedagogically superior to the standard phrasing is a claim about teaching, not about mathematics, and this chapter does not have the evidence to settle it; it is offered as a candidate, not a result.*

5.4 What This Chapter Does Not Yet Give

The honest limitation is more important than either of the two positive results above, and stating it clearly is what keeps this chapter from overclaiming on behalf of the rest of the book.

Observation 11. *The amplitwist reading used throughout this author’s work on locomotion depended on a population: many coupled oscillators, an order parameter $Z = re^{i\psi}$ measuring their coherence, and a hierarchy of such populations composed across scales. Nothing of that kind is present in a single two-terminal impedance. A single $Z(j\omega)$ has an amplitude and a phase because it is a single complex number, not because a population of anything has cohered around a common phase. The richer apparatus — coherence, order parameters, hierarchical composition of coupled populations — has no referent yet at the*

level of a single impedance, and will not acquire one until this book reaches oscillator populations, synchronization, and the phase-locked loop. This chapter's amplitwist is Needham's original, single-point notion; it is not yet the Kuramoto-flavored population notion this author's other work builds on, and the two should not be conflated before the book has earned the second one.

5.5 Verdict

The chapter closes by answering the questions it was written to answer, rather than assuming a favorable answer along the way.

Does calling impedance an amplitwist operator generate a new quantitative prediction beyond ordinary phasor analysis? No. Every equation in this chapter is standard circuit theory; nothing here computes a number an electrical engineer could not already compute.

Does it generate a genuine conceptual clarification an electrical engineer would recognize as useful rather than merely relabeled? Partially, and in one specific place: the forced distinction between operator addition (series and parallel combination) and operator composition (cascading, reserved for Chapter 6) is a real, if modest, sharpening of something ordinary pedagogy leaves as two unconnected memorized rules. The resonance reframing is a candidate for pedagogical value, honestly stated as unproven rather than as a demonstrated result.

Is the pilot's verdict sufficient to justify the rest of the book? Conditionally. This chapter does not establish that amplitwist is the natural language of circuit theory in general; it establishes that impedance is amplitwist trivially, and that the vocabulary earns a small, real dividend in conceptual clarity at this level while producing no new predictions. The larger claim this book is built on — that circuits are a continuation system in the stronger sense developed for locomotion — depends on the population-level chapters (two-port cascades, synchronization, oscillators, phase-locked loops) doing more work than this one did, in the same way the Kuramoto order parameter, not the bare notion of a phasor, carried the weight in the locomotion essay. Chapter 5 has earned the right to a Chapter 6. It has not, by itself, earned the whole book.

Chapter 5 Quick Reference

- * Impedance $Z = |Z|e^{j\theta}$ acting on current phasor I to produce voltage phasor V is, by definition, an amplitwist operator (amplitude $|Z|$, twist θ). This identification is exact, not interpretive.
- * Series/parallel combination is **operator addition** ($Z_1 + Z_2$ or $Y_1 + Y_2$), applied to a shared input. Cascading circuit stages is **operator composition** (Chapter 6), applied sequentially. These are different operations and are often taught as unrelated rules.
- * Resonance is the condition where twist $\theta \rightarrow 0$: the operator degenerates from rotation-and-scaling to pure scaling. Standard result, reframed geometrically.
- * This chapter's amplitwist has no population and no coherence measure. The Kuramoto-style order-parameter reading (population coherence, hierarchical composition) is not yet available at the single-impedance level; it begins in the synchronization and oscillator chapters.
- * **Verdict:** No new predictions at this level; one real conceptual clarification (addition vs. composition); the book's larger claim is not yet earned by this chapter alone.

6 Two-Port Networks as Amplitwist Chains

Chapter 5 ended by deferring a promise: series and parallel combination of impedances are additive operations on operators sharing a common input, and the genuinely compositional operation this book's framework depends on — one operator's output becoming the next operator's input — was reserved for cascaded two-port networks. This chapter cashes that promise, and holds itself to the same standard Chapter 5 set: state plainly what is exact, what is a genuine result, and what is a boundary condition beyond which the scalar amplitwist picture stops applying.

6.1 Cascading Is Matrix Composition, Not Scalar Composition

A two-port network relates the voltage and current at its input port to the voltage and current at its output port through the transmission (ABCD) parameters,

$$\begin{pmatrix} V_1 \\ I_1 \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} V_2 \\ -I_2 \end{pmatrix}.$$

When two such networks are cascaded, the output port of the first connected directly to the input port of the second, the combined network's ABCD matrix is the

ordinary matrix product of the two individual matrices,

$$M_{\text{total}} = M_1 M_2.$$

This is genuine sequential composition, in exactly the sense Chapter 5 deferred: the output of the first stage's transformation is the literal input to the second stage's transformation.

Observation 12. *It would overreach to call this, without qualification, an “amplitwist chain” in the sense of Chapter 5. A single amplitwist operator, as Needham defines it and as Chapter 5 uses it, is a single complex number acting by multiplication: a map from \mathbb{C} to \mathbb{C} . An ABCD matrix is a linear map on the two-dimensional space of port variables (V, I) ; it is a genuine composition, but of a more general kind of operator than a scalar amplitwist. Chapter 5's amplitude-and-twist decomposition does not automatically survive the move from a scalar multiplier to a 2×2 matrix, and this chapter should not claim it does without showing where, specifically, it does.*

6.2 Where the Scalar Picture Is Recovered Exactly

The scalar amplitwist picture returns, exactly rather than approximately, in a specific and identifiable case: a cascade of matched, reflectionless two-port stages, characterized by a single complex transmission coefficient $T = |T| e^{j\phi}$ (the forward transmission scattering parameter, familiar from microwave and RF circuit theory) rather than by a full ABCD matrix.

Definition 6 (Matched cascade). *A two-port stage is matched at a given reference impedance if it presents no reflection at either port when terminated in that reference impedance, so that its behavior is fully captured by a single complex transmission coefficient $T = |T| e^{j\phi}$, with $|T|$ the magnitude (gain or attenuation) and ϕ the insertion phase.*

Proposition 4 (Transmission coefficients compose like amplitwist operators). *Let $T_1 = |T_1| e^{j\phi_1}$ and $T_2 = |T_2| e^{j\phi_2}$ be the transmission coefficients of two matched two-port stages, connected in cascade, both referenced to the same characteristic impedance so that no reflection occurs at the junction between them. Then the transmission coefficient of the cascade is*

$$T_{\text{total}} = T_1 T_2 = |T_1| |T_2| e^{j(\phi_1 + \phi_2)}.$$

Proof. Under the matched condition, the signal leaving stage 1 is not partially reflected back into stage 1 by stage 2's input, so the output of stage 1 is transmitted into stage 2 exactly as it would be into a matched load; the overall transmission is the product of the two stages' individual transmission coefficients, with no additional interference terms. This is the standard result for cascaded matched networks in microwave circuit theory (see, e.g., Pozar's treatment of cascaded two-port networks via scattering parameters). \square

Observation 13. *This is not a metaphorical use of amplitwist composition; it is the identical algebraic fact that governs the composition of holomorphic maps. If f and g are holomorphic and $h = f \circ g$, the chain rule gives $h'(z) = f'(g(z)) g'(z)$: derivatives compose by multiplication, magnitudes multiply and arguments add. Proposition 2 is this same rule, applied to transmission coefficients rather than to derivatives of a map. A chain of matched two-port stages is, in the precise sense Chapter 5 was careful not to claim for the general ABCD case, literally an amplitwist chain: amplitude multiplies, twist (phase) adds, at every stage.*

Two concrete cases make the two components of this composition visible separately. A chain of n identical resistive attenuators, each with real transmission coefficient $|T_i| < 1$ and negligible phase shift, composes to $T_{\text{total}} = \prod_i |T_i|$: pure amplitude composition, no twist, which is why cascaded attenuator losses are added in decibels (logarithms of the multiplied magnitudes) in ordinary engineering practice. A chain of n matched transmission-line sections of electrical length ϕ_i and negligible loss composes to $T_{\text{total}} = e^{j \sum_i \phi_i}$: pure twist composition, no amplitude change, which is why cascaded transmission-line phase shifts simply add. An amplifier chain with both gain and phase shift at each stage is the general case, combining both.

6.3 Assembly Index of a Two-Port Cascade

Definition 7 (Assembly Index of a matched cascade). *For a cascade of n matched two-port stages composed as in Proposition 2, define the Assembly Index of the cascade as $A(\text{cascade}) = n - 1$, the number of compositions (multiplications of transmission coefficients) required to reach the overall T_{total} from the n individual stage coefficients.*

This is the direct analogue, for genuine operator composition, of the Assembly Index defined for the additive amplitwist chain of the locomotion volume, and the two should be kept distinct: that Assembly Index counted applications of a coupling functional \mathcal{C} combining a population into a coarser description; this one counts applications of ordinary complex multiplication combining sequential stages into a cascade. Both are legitimate causal-depth measures, but they are depth measures for different kinds of composition, additive-hierarchical in one case and sequential-multiplicative in the other, and this book will need both before it is finished.

6.4 Where the Scalar Picture Breaks Down

Observation 14. *The matched-cascade result of Proposition 2 is exact only under the reflectionless assumption stated in Definition 1. When two cascaded stages are mismatched — when a stage’s output impedance does not equal the next stage’s input reference impedance — reflections at the junction produce multiple internal reflection terms between the stages,*

and the overall transmission coefficient of the cascade is no longer simply the product T_1T_2 . The correct general treatment returns to the full linear-operator picture of Section 6.1: ABCD matrices (or, equivalently, chain scattering parameters constructed for exactly this purpose) multiply as matrices under cascading, capturing the reflection interference terms that scalar multiplication of transmission coefficients cannot. The scalar amplitwist chain of Section 6.2 is therefore a special, idealized case of the general composition law, exact precisely at impedance matching and increasingly inexact as mismatch grows, in the same way Chapter 3 of the locomotion volume noted that the population-level Kuramoto order parameter is a mean-field idealization rather than an exact microscopic description.

This boundary is worth stating with the same explicitness given to the ground-contact discontinuity in the locomotion volume's failure-mode chapter: mismatch is not a rare edge case in circuit design, it is the normal condition whenever stages are not deliberately impedance-matched, and a reader should not come away from Section 6.2 believing scalar amplitwist composition is the general story. It is the matched limit of a more general, matrix-valued composition law.

6.5 Verdict

Does this chapter's amplitwist reading generate something beyond restatement? Yes, and more substantially than Chapter 5. Proposition 2 is not a relabeling; it identifies cascaded matched-network transmission as a literal instance of the same algebraic law that governs composition of holomorphic derivatives, and derives from it the ordinary engineering practice of adding decibels and adding phases across a matched chain, recovering a familiar practical rule as a special case of amplitwist composition rather than as an independent piece of engineering folklore.

Does the chapter also state where this breaks down? Yes: mismatch is the general case, not the exception, and Section 6.5 says so without qualification. The scalar amplitwist chain of this chapter is earned exactly at the matched limit and requires the full matrix composition of Section 6.1 otherwise.

Chapter 5 earned a conditional pass: real but modest clarification, no new predictions. This chapter earns a stronger pass: a genuine derived result (amplitude multiplies, phase adds, exactly the chain rule for composed holomorphic maps) with an explicitly bounded domain of exactness. The book's larger claim is still not settled by two chapters, but the two together now show both halves of the amplitwist decomposition — addition without composition in Chapter 5, composition without addition in Chapter 6 — doing distinguishable, non-overlapping, non-trivial work.

7 Resonance as an Amplitwist Fixed Point

Chapter 5 introduced impedance as an amplitwist operator and noted, briefly, that resonance may be understood as the condition under which the operator's twist vanishes. Chapter 15 later returned to the same observation in the context of oscillator circuits, where resonance removed the phase contribution of the feedback network and left the active device to satisfy the Barkhausen phase condition alone. Both appearances were deliberately brief. This chapter develops the observation in its own right.

The goal is not to replace the standard theory of resonance. The goal is to determine whether the amplitwist vocabulary identifies a structural feature that ordinary circuit language leaves implicit.

7.1 Resonance in the Standard Vocabulary

Consider a series RLC circuit with impedance

$$Z(j\omega) = R + j \left(\omega L - \frac{1}{\omega C} \right).$$

The resonance frequency is defined by the condition

$$\omega_0 L = \frac{1}{\omega_0 C},$$

which yields

$$\omega_0 = \frac{1}{\sqrt{LC}}.$$

At this frequency,

$$Z(j\omega_0) = R,$$

a purely real impedance. This result is elementary and appears in every introductory treatment of resonant circuits.

Observation 15. *Nothing in this chapter disputes the standard interpretation. At resonance the inductive and capacitive reactances cancel, leaving only the resistive component. The question is whether the amplitwist interpretation reveals an additional structural feature of the same fact.*

7.2 The Twist of an Impedance Operator

Chapter 5 defined impedance in polar form, $Z(j\omega) = |Z(j\omega)|e^{j\theta(\omega)}$, and identified the magnitude and phase as amplitude and twist respectively.

Definition 8 (Twist of an impedance). *The twist of an impedance operator is the phase angle $\theta(\omega) = \arg Z(j\omega)$. Equivalently, it is the angle through which the impedance rotates the current phasor in producing the voltage phasor.*

For the series RLC impedance,

$$\theta(\omega) = \tan^{-1} \left(\frac{\omega L - \frac{1}{\omega C}}{R} \right).$$

At resonance the numerator vanishes, so $\theta(\omega_0) = 0$.

Proposition 5. *A series RLC circuit is resonant if and only if its impedance amplitwist operator has zero twist.*

Proof. By definition, $\theta(\omega) = \tan^{-1} \left(\frac{\omega L - \frac{1}{\omega C}}{R} \right)$. The twist vanishes exactly when $\omega L - \frac{1}{\omega C} = 0$, which is precisely the resonance condition. \square

This proposition contains no new mathematics. Its significance lies elsewhere. Ordinary circuit language states that reactances cancel. The amplitwist language states that the operator stops rotating. These are equivalent descriptions of the same event.

7.3 Fixed Points of Operator Action

The more interesting question concerns why resonance appears repeatedly throughout circuit theory rather than only in isolated examples. The amplitwist interpretation suggests an answer. Away from resonance, the impedance operator performs two actions simultaneously: scaling and rotation. At resonance, one of these actions disappears; only scaling remains.

Definition 9 (Amplitwist fixed point). *Let $A(\omega) = r(\omega)e^{j\theta(\omega)}$ be an amplitwist operator parametrized by frequency. A frequency ω_0 is an amplitwist fixed point if $\theta(\omega_0) = 0$, so that the operator reduces to pure scaling.*

Observation 16. *The phrase “fixed point” is being used geometrically here, and should not be confused with the dynamical, self-referential sense of the same phrase used for feedback loops in Chapter 13 (where X_o was a fixed point of the map $g(X_o) = A(X_i - \beta X_o)$). The claim here is not that any trajectory ceases evolving or that any equation is being solved for a self-consistent value; it is only that the rotational component of a frequency-dependent operator vanishes at a particular frequency, leaving pure amplitude transformation. The two uses of “fixed point” in this book are related only by loose family resemblance (both describe a special, structurally simpler condition within a larger space of possibilities) and should be read independently of one another.*

Resonance is therefore the simplest example of an amplitwist fixed point in this geometric sense: the operator ceases to rotate and acts only by scaling.

7.4 Resonance and Admissibility

Chapter 14, later in this book, interprets oscillation as an admissibility condition. The Barkhausen criterion requires the total loop phase shift to vanish modulo a full rotation. Resonance can now be read as a local version of the same requirement, previewed briefly here and developed formally there.

Proposition 6. *For a resonant network embedded within a feedback loop, resonance minimizes the phase contribution of the network to the loop's total admissibility condition.*

Proof. At resonance the network's twist vanishes. The phase contribution of the network to the total loop phase is therefore zero. Consequently the burden of satisfying the Barkhausen phase condition is reduced to whatever phase contributions remain elsewhere in the loop. \square

Observation 17. *This is why resonant tanks appear so naturally in oscillator design. The tank is not merely selecting a preferred frequency. It is selecting a frequency at which its own rotational action disappears. The feedback loop therefore encounters the least phase opposition to self-sustained continuation at precisely that frequency.*

7.5 Resonance and Synchronization

The Synchronization chapter, later in this book, introduces phase locking as the emergence of coherence among coupled oscillators. Resonance is relevant there as well, previewed briefly here: an oscillator driven near resonance requires less phase correction from an external signal than one driven far from resonance.

Observation 18. *A driven oscillator operating near resonance is already close to a zero-twist condition. The external coupling therefore spends less effort compensating rotational mismatch and more effort maintaining coherence. Resonance and synchronization are not identical phenomena, and this connection is stated here only as a qualitative observation, not as a derived proposition; a fully rigorous treatment would need to relate the impedance twist of a driven resonant circuit to the phase-error variable of Adler's equation explicitly, which this chapter does not attempt.*

This observation helps explain why resonant structures appear repeatedly in synchronized systems, from oscillator arrays to phase-locked loops, while stopping short of claiming a proven equivalence between the two phase concepts involved.

7.6 Where the Fixed-Point Picture Breaks Down

The amplitwist interpretation has limits. Not every resonant phenomenon reduces cleanly to the vanishing of a scalar twist.

Observation 19. *Distributed systems such as transmission lines, cavity resonators, and multimode structures may support several simultaneous resonances, each with its own phase structure. The scalar amplitwist picture developed here is exact only when a single effective impedance operator adequately describes the system.*

Likewise, strongly nonlinear oscillators can exhibit amplitude-dependent resonance frequencies, mode hopping, and chaotic behavior that resist description as a single fixed point of a frequency-dependent operator. These are not defects of the amplitwist language; they are indications of where the scalar picture must give way to a higher-dimensional one.

7.7 Verdict

Chapter 5 asked whether the amplitwist vocabulary could do real conceptual work rather than merely rename existing concepts. Resonance provides one of the clearer affirmative examples. The standard description says that reactances cancel; the amplitwist description says that rotational action vanishes. The two statements are mathematically equivalent, but the second names a structural pattern (minimal twist) that recurs, in qualitatively similar but not identically proven form, across admissibility, synchronization, and oscillation elsewhere in this book. Resonance is, in the precise and narrow sense established here, the simplest amplitwist fixed point available in elementary circuit theory.

8 Filters as Admissibility Gates in Frequency Space

Chapter 7 treated resonance as an amplitwist fixed point: a frequency at which an operator's twist vanishes and only scaling remains. That chapter focused on a single frequency. This chapter widens the lens.

Real circuits rarely care about one frequency alone. They must decide which frequencies are permitted to continue through a system, which are attenuated, and which are rejected altogether. Ordinary circuit theory describes this behavior in terms of transfer functions, passbands, stopbands, and cutoff frequencies. This chapter asks whether these familiar structures admit a more general reading as admissibility conditions operating over frequency space.

The claim is not that filters secretly contain a new mathematics. The claim is that filtering may be the clearest circuit-theoretic realization of admissibility yet encountered in this book.

8.1 A Filter as a Selection Operator

Let $H(j\omega) = |H(j\omega)|e^{j\theta(\omega)}$ denote the transfer function of a linear time-invariant filter. For an input spectrum $X(j\omega)$, the output spectrum is $Y(j\omega) = H(j\omega)X(j\omega)$.

Observation 20. *This equation already has the form encountered throughout Chapters 5 through 7. The transfer function is an amplitude operator acting on a spectral input. The distinction is that the operator is now evaluated across an entire frequency domain rather than at a single frequency.*

A filter therefore acts not on a signal as a whole but on the signal's frequency components individually. Some components survive. Others do not.

8.2 Admissibility in Frequency Space

The language of admissibility becomes natural once the transfer function is viewed as a selection mechanism.

Definition 10 (Frequency admissibility). *Let $A = \{\omega : |H(j\omega)| \geq \tau\}$ for some threshold τ . A frequency is admissible if it belongs to A and inadmissible otherwise.*

Observation 21. *The threshold τ is not unique. Different engineering applications tolerate different levels of attenuation. The admissibility concept therefore depends on context, exactly as admissibility does elsewhere in this author's work.*

The important point is structural. A filter partitions frequency space into two regions. One region continues. The other does not.

8.3 Low-Pass Filters as Continuation Preservers

Definition 11 (Ideal low-pass admissibility set). *An ideal low-pass filter with cutoff frequency ω_c defines $A = \{\omega : |\omega| \leq \omega_c\}$. Frequencies inside A are admissible; frequencies outside A are refused.*

Observation 22. *The usual engineering description is that the filter passes low frequencies and attenuates high frequencies. The admissibility description says the same thing in different language: low frequencies satisfy the continuation criterion imposed by the filter while high frequencies do not. Nothing new has been discovered; a structural distinction has merely been made explicit.*

8.4 Band-Pass Filters and Admissibility Windows

Definition 12 (Admissibility window). *A band-pass filter defines an admissible region $A = \{\omega : \omega_1 \leq |\omega| \leq \omega_2\}$. The admissible set is neither everything below a threshold nor everything above one, but a bounded interval.*

Observation 23. *Many physical systems operate exactly this way. A radio receiver does not ask whether a signal exists; it asks whether the signal lies within a narrow admissible region centered on the desired carrier frequency. The filter is therefore functioning as an*

admissibility gate long before any digital decoding occurs. This pattern is independent of electronics: it appears whenever a system permits only a restricted region of a larger possibility space.

8.5 Stopbands as Refusal Regions

This book's Chapter 9 will introduce, through the diode, a Refuse operator: a mapping that collapses every inadmissible input to an identical null output. Filters provide an earlier and independent realization of a structurally similar pattern, in the frequency domain rather than the voltage domain.

Definition 13 (Frequency refusal region). *Let $R = \{\omega : |H(j\omega)| < \tau\}$. The set R is the refusal region of the filter.*

Observation 24. *A diode (Chapter 9) refuses according to voltage polarity; a filter refuses according to frequency. The selection criterion differs, but the structural role, collapsing an entire inadmissible region to a uniformly rejected output, anticipates the same pattern Chapter 9 will treat formally, without depending on that chapter's content here. The filter occupies an intermediate position between the purely reversible amplitwist operators of Chapters 5 through 7 and the explicitly irreversible operator Chapter 9 introduces.*

8.6 Bandwidth as Admissibility Volume

Definition 14 (Admissibility volume). *The admissibility volume of a filter is the measure of its admissible frequency region, $V_A = \int_A d\omega$. For a simple band-pass filter this reduces to its bandwidth.*

Proposition 7. *Increasing bandwidth increases admissibility volume; decreasing bandwidth decreases it.*

Proof. Immediate from Definition 4. □

Observation 25. *This is not a profound theorem. Its importance lies in the vocabulary it introduces: bandwidth can now be discussed as a volume of admissible spectral behavior rather than merely as a numerical interval. This language becomes useful later, when this book discusses information, noise, and signal-to-noise ratio.*

8.7 Transition Regions and Boundary Effects

Ideal filters possess sharp boundaries. Real filters do not.

Observation 26. *Every realizable filter contains a transition region between passband and stopband. The filter does not draw a perfectly sharp line between admissible and inadmissible frequencies; it contains a boundary layer in which admissibility becomes progressively less certain. The transition region is neither fully accepted nor fully rejected; it occupies an intermediate status.*

Observation 27. *The sharper the filter, the steeper this boundary becomes. Infinite sharpness corresponds to an idealized discontinuity that physical circuits cannot realize exactly.*

8.8 Resonance and Filtering

Proposition 8. *Every resonant filter defines an admissibility region centered on an amplitwist fixed point.*

Proof. Resonance (Chapter 7) selects a characteristic frequency ω_0 at which the network's impedance twist vanishes. The transfer function magnitude is maximized in a neighborhood of ω_0 . This neighborhood therefore forms the center of the filter's admissible region. \square

Observation 28. *Resonance identifies where continuation is easiest. Filtering determines how far from that point continuation remains acceptable. The two concepts are complementary rather than independent.*

8.9 Where the Admissibility Picture Becomes Incomplete

Observation 29. *Not every filter can be reduced cleanly to a binary admissibility gate. Adaptive filters, nonlinear filters, and time-varying filters may alter their admissibility regions dynamically in response to incoming signals; the admissibility region then becomes a moving target rather than a fixed subset of frequency space. Filters operating on multi-dimensional signal spaces may require admissibility criteria that cannot be represented by a single frequency interval. The present chapter is best understood as a foundation rather than a complete theory.*

8.10 Verdict

Resonance identified a single frequency at which an amplitwist operator becomes purely scaling. Filtering generalizes that observation from a point to a region. A filter determines which frequencies may continue through a system and which may not: passbands become admissible regions, stopbands become refusal regions anticipating Chapter 9's formal treatment of refusal, bandwidth becomes admissibility volume, and transition regions become boundary effects. None of these reinterpretations change the mathematics of filtering. What they provide is a common language linking circuit theory to the broader admissibility framework developed elsewhere in this book.

9 The Diode as an Irreversible Spherpoper Operator

Chapters 5 and 6 developed two circuit operators, impedance and the ABCD matrix, that are both, in the relevant sense, reversible: given the output of either opera-

tor and the operator itself, the input can always be recovered (Chapter 5's $I = V/Z$ for $Z \neq 0$; Chapter 6's port-2 variables from port-1 variables and M^{-1} , wherever $\det M \neq 0$). This chapter introduces the first circuit element in this book for which that is no longer true, and asks what changes when the operator vocabulary must include genuine, unrecoverable collapse rather than scaling, rotation, and their composition.

9.1 Recalling the Pop Operator

This author's prior work on procedural ontology defines a local collapse operator for bounded evaluable regions: given a syntactically closed region E , delimited unambiguously from its surroundings, the Pop operator P maps E to a value $E' = P(E)$, substituting that value for the region it occupied, using no information from outside E 's own boundary. The definition is deliberately general, covering arithmetic evaluation, geometric transformation, and attribute assignment as instances of a single underlying pattern: locate a bounded, well-formed region, collapse it, splice the result back in.

Observation 30. *The Pop operator, as originally defined, does not specify whether the collapse is invertible. Arithmetic evaluation is generally not invertible ($3+4$ pops to 7; nothing in the value 7 recovers the specific addends). This chapter's contribution is to ask what it means, concretely, for a physical circuit element to perform a Pop-like collapse, and to identify the diode as the clearest such case in elementary circuit theory.*

9.2 The Ideal Diode as Pop Paired with Refuse

An ideal diode conducts, ideally without resistance, when forward biased, and blocks all current, ideally regardless of the reverse voltage's magnitude, when reverse biased. This chapter introduces a second operator, dual to Pop, to describe the blocking behavior; unlike the Pop operator above, this operator is proposed here rather than recovered from prior formal work, and should be read as a new contribution to this author's Spherepop vocabulary rather than a restatement of one.

Definition 15 (Refuse operator, introduced here). *Let E be a bounded region presented as a candidate for collapse. The Refuse operator \bar{P} maps E to a fixed null output $\bar{P}(E) = \emptyset$ whenever E fails an admissibility test, without regard to any further distinguishing detail of E beyond its failing that test. Where Pop collapses an admissible region to a value that depends on the region's specific content, Refuse collapses an inadmissible region to a value that does not.*

Definition 16 (Ideal diode as a Pop/Refuse pair). *Let V be the voltage across an ideal*

diode. Define the admissibility test $V > 0$ (forward bias). The diode's action is

$$D(V) = \begin{cases} P(V) = V/R_{on} & V > 0 \quad (\text{admissible: Pop, forward conduction}) \\ \bar{P}(V) = 0 & V \leq 0 \quad (\text{inadmissible: Refuse, reverse blocking}). \end{cases}$$

Observation 31. Every reverse voltage, from $-\epsilon$ to $-\infty$, is mapped by \bar{P} to the identical output 0. This is precisely the information-collapsing character the Pop/Refuse vocabulary is meant to name: the specific magnitude of the reverse voltage is not preserved anywhere in the output, and no further operation on the output alone can recover it. This is a stronger claim than “the diode is nonlinear”; resistors driven into a nonlinear regime can still be invertible. The diode's reverse branch is not merely nonlinear, it is many-to-one, and many-to-one is what makes an operator irreversible in the sense this chapter needs.

9.3 Irreversibility Is Not an Incidental Feature: Rectification Requires It

The strongest claim this chapter can make is not that the diode happens to be well-described by an irreversible operator, but that performing its characteristic function, rectification, is impossible for any invertible operator whatsoever. This is worth proving rather than asserting, since it converts a descriptive choice into a necessity.

Theorem 1 (Rectification requires a non-invertible operator). *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be any single-valued, invertible (injective) function. Then $f(V) \neq f(-V)$ for every $V \neq 0$. Consequently, no invertible circuit operator can perform rectification, understood as producing an output that depends only on $|V|$ (or otherwise collapses the sign of V), since such an output is, by definition, identical for V and $-V$.*

Proof. If f is injective, $f(a) = f(b)$ implies $a = b$. For any $V \neq 0$, $V \neq -V$, so injectivity forces $f(V) \neq f(-V)$. An operator that rectifies must satisfy $f(V) = f(-V)$ for at least the pairs it rectifies (by definition of collapsing the sign), which directly contradicts injectivity for any such pair. Hence no injective, and therefore no invertible, function can rectify. \square

Corollary 1. *Since Chapters 5 and 6 established that impedance and cascaded two-port networks (in the ideal, matched, or generally invertible case) are reversible operators, neither an arrangement of ideal resistors, inductors, capacitors, nor matched transmission stages alone can rectify a signal. Rectification is not merely a convenient application of diodes; it is a task that provably requires an operator outside the amplitwist vocabulary of Chapters 5 and 6 altogether. The Pop/Refuse pair of Definition 2 is not an alternative description of what a diode does; it names the only kind of operator that could do it.*

This is the chapter's strongest result, and it is worth being clear about what kind of result it is. Theorem 1 is elementary — a direct consequence of the definition of

injectivity — but its elementariness is exactly what makes it valuable here: it is not a claim that depends on choosing to describe diodes in Spherepop vocabulary, it is a claim about any conceivable rectifying operator, diode-based or otherwise, and it is true independent of anything else in this book.

9.4 Dissipation: Where Entropy Enters This Book

Real diodes are not ideal switches; forward conduction follows the exponential diode law $I = I_s (e^{V/\eta V_T} - 1)$, and forward conduction dissipates real power, $P = VI > 0$, as heat. This is the first place in this book where the entropy field S of the RSVP triad (Φ, v, S) becomes unavoidable rather than a background assumption: Chapters 5 and 6 could, in the ideal lossless limit, be treated as conservative, reversible transformations, but a conducting diode cannot, even ideally, since forward conduction is inherently dissipative in a way that ideal inductors and capacitors are not.

Observation 32. *It is tempting to reach further and invoke Landauer’s bound directly, treating each act of reverse-bias blocking as an erasure of one bit of sign information at a definite thermodynamic cost of $k_B T \ln 2$. This chapter declines to make that claim. Landauer’s bound concerns discrete, clocked logical erasure; a diode is a continuous, unclocked nonlinear device, and asserting a literal per-erasure thermodynamic cost for continuous rectification would overreach what is actually established. What can be said honestly is weaker and still substantial: forward conduction is unavoidably dissipative, this dissipation is a genuine instance of entropy production rather than a modeling inconvenience, and the connection to Landauer-style reasoning about information erasure is a suggestive open question for this book’s later chapters on noise and entropy, not a result this chapter is entitled to claim.*

9.5 Where the Two-State Model Breaks Down

The Pop/Refuse pair of Definition 2 is a considerable idealization, and this book’s discipline requires stating plainly where it fails.

Observation 33. *Three failures are worth naming explicitly. First, real diodes exhibit reverse breakdown: beyond a sufficiently negative voltage (the Zener or avalanche breakdown voltage), the diode conducts in reverse, which the two-branch model of Definition 2 has no room for and which requires a third regime entirely. Second, and more importantly for this book’s larger themes, real diode switching exhibits reverse recovery: a diode conducting in the forward direction, upon having its bias suddenly reversed, continues to conduct briefly in the reverse direction while stored minority-carrier charge in the junction is removed. This is not captured by any instantaneous function of the present voltage $V(t)$ alone; it depends on the diode’s recent conduction history, in exactly the sense this book’s planned chapter on histories before states will make precise for capacitors and inductors. The diode, in this*

respect, is not only irreversible in the collapse sense of Definition 2, it is also historical in the stronger sense of Chapter 5.4 of this author's work on locomotion: its present behavior is a function of accumulated past state, not of an instantaneous input alone. Third, junction capacitance means the transition between forward and reverse states is not instantaneous even apart from stored charge, introducing genuine continuous dynamics the discrete Pop/Refuse model omits entirely.

None of these three failures undermines Theorem 1, which concerns only the idealized rectifying function and holds regardless of how any particular device approximates it. They do mean that Definition 2's two-branch model is a first approximation, useful for establishing why irreversibility is necessary at all, but not a complete account of what a real diode does across time.

9.6 Verdict

Does this chapter's operator reading generate something beyond restatement? Yes, and its central result is a proof rather than a redescription: Theorem 1 establishes that rectification requires a non-invertible operator as a matter of necessity, independent of whether the reader accepts any of this book's other vocabulary. This is a stronger and more self-standing result than either Chapter 5's modest clarification or Chapter 6's derived composition law, precisely because it does not depend on accepting the amplitwist framework to be true; it depends only on the definition of an invertible function.

Does the chapter state its own boundaries honestly? Yes: the Landauer connection is explicitly declined rather than asserted, and the three failure modes of Section 9.5 are named rather than smoothed over, including an early, honest pointer toward the history-dependence (reverse recovery) that this book's Histories-before-States material will need to treat properly rather than in passing.

Chapter 9 Quick Reference

- * Chapters 5–6 dealt only with **reversible** (invertible) operators. The diode is this book’s first **irreversible** operator.
- * Ideal diode = Pop (forward, $V > 0$: collapses to V/R_{on}) paired with **Refuse** (reverse, $V \leq 0$: collapses every input to the same null output 0). Refuse is introduced in this chapter, not recovered from prior work.
- * **Theorem:** No invertible operator can rectify (collapse sign information), since injectivity forbids $f(V) = f(-V)$ for $V \neq 0$. Irreversibility is therefore *necessary* for rectification, not merely descriptive.
- * Forward conduction dissipates real power ($P = VI > 0$): this is where the entropy field S becomes unavoidable in this book. The tempting Landauer (bit-erasure) connection is explicitly **declined** as unrigorous for a continuous device — flagged as an open question, not claimed as a result.
- * **Breaks down:** reverse breakdown (needs a third regime); reverse recovery (diode behavior depends on *recent conduction history*, not just instantaneous $V(t)$ — an early link to the Histories-before-States chapter); junction capacitance (continuous dynamics the discrete model omits).
- * **Verdict:** Strongest chapter yet — a genuine, self-standing proof, independent of whether the reader accepts the rest of the book’s framework.

Histories Before States

Chapter 9 noted that diode reverse recovery depends on recently stored junction charge, not on instantaneous voltage alone, and pointed here rather than treating it in passing. This chapter now supplies that foundation, and it does so by starting from the two circuit elements for which the claim is not an interpretation at all: the capacitor and the inductor. Two later chapters will draw directly on what is established here: Chapter 13’s treatment of feedback paths containing reactive elements, and the Phase-Locked Loop chapter’s account of the loop filter’s integrator as a genuine historical state, structurally unlike anything available to a bare injection-locked oscillator. Both are previewed briefly later in this chapter, in Section 1.3 below, so that this chapter’s foundation is visibly connected to where it will be used, without requiring those later chapters’ own content in advance.

The Capacitor and Inductor Are, Literally, Integrals Over History

Definition 17 (Capacitor and inductor constitutive relations, integral form). For a capacitor with capacitance C and an inductor with inductance L ,

$$v_C(t) = v_C(t_0) + \frac{1}{C} \int_{t_0}^t i(\tau) d\tau, \quad i_L(t) = i_L(t_0) + \frac{1}{L} \int_{t_0}^t v(\tau) d\tau.$$

Observation 34. Standard circuit pedagogy usually presents these relations in differential form, $i = C\dot{v}$ and $v = L\dot{i}$, which is mathematically equivalent but rhetorically obscures exactly the point this chapter needs. The differential form makes a capacitor look like a device whose current depends on the instantaneous rate of change of its voltage; the integral form makes visible what is actually true, that the voltage at any moment is nothing but the accumulated record of every current that has ever flowed through the device since some reference time. A capacitor is not a device with a voltage that happens to relate to current by a derivative; it is, in the most literal sense available anywhere in elementary circuit theory, a device defined by its own history.

The State Is a Lossy Compression of the History

This chapter's central claim is stronger than the observation above, and is worth proving rather than asserting: the value $v_C(t_0)$ used as an "initial condition" is not merely derived from history, it actively discards information the history contained.

Proposition 9 (Capacitor voltage is a non-injective functional of current history). Let $i_1(\tau)$ and $i_2(\tau)$ be two distinct current waveforms on $(-\infty, t_0]$ (i.e. $i_1 \neq i_2$ as functions) with equal total delivered charge, $\int_{-\infty}^{t_0} i_1(\tau) d\tau = \int_{-\infty}^{t_0} i_2(\tau) d\tau$. Then the resulting capacitor voltage $v_C(t_0)$ is identical under either history.

Proof. Immediate from Definition 1: $v_C(t_0)$ depends on the current history only through the value of its integral, $\frac{1}{C} \int_{-\infty}^{t_0} i(\tau) d\tau$. Since integration is a linear functional, any two current histories with the same integral produce the same $v_C(t_0)$, regardless of how differently the current was distributed across time to produce that total. Concretely, a slow, steady trickle of charge and a sharp early pulse followed by silence can deliver identical total charge and are therefore indistinguishable in $v_C(t_0)$, even though they are different functions. \square

Corollary 2. The capacitor voltage $v_C(t_0)$ is, in the precise sense developed throughout this author's other work, a state: a compressed summary of a trajectory (the current history) that discards the path by which it was reached in exchange for a smaller, tractable object. Proposition 1 makes this provable rather than assertable for a standard textbook circuit element: infinitely many distinct current histories collapse to the same state, and only the state, not the discarded history, is available to determine the capacitor's future behavior given future inputs.

Observation 35. *This is also, without needing to force the connection, the precise technical content behind standard circuit and control theory’s own terminology. Capacitor voltages and inductor currents are conventionally called state variables in exactly the sense this book has been using the word: the minimal set of quantities such that, together with future inputs, they suffice to determine all future behavior, with everything else about the past rendered irrelevant. Ordinary engineering vocabulary already agrees with this chapter’s claim; it is not importing a foreign framework, it is making explicit why the word “state” was the right word to begin with.*

Making Good on Three Debts

Observation 36 (Diode reverse recovery, promised in Chapter 9). *The stored minority-carrier charge responsible for diode reverse recovery, conventionally written Q_{rr} , is itself a history-integral of exactly the kind formalized in Definition 1: Q_{rr} is (to good approximation) the integral of recent forward current, $Q_{rr} \approx \int i_F(\tau) d\tau$ over the recent conduction interval. A diode near a switching transition is, in this respect, behaving as a capacitor-like history-integrating element layered on top of the Pop/Refuse collapse of Chapter 9, not as an alternative to it: the same device is non-invertible in the instantaneous sense of Chapter 9 (many reverse voltages collapse to the same zero output) and historical in the integrative sense of this chapter (its transition behavior depends on accumulated recent forward conduction). These are compatible, not competing, descriptions of two different aspects of the same physical device.*

Observation 37 (The PLL loop-filter integrator, used by the Phase-Locked Loop chapter later in this book). *The Type-II loop filter’s integrator output, responsible for the zero-steady-state-error result proven there, is exactly $\frac{1}{\tau_I} \int V_d(\tau) d\tau$: an accumulated integral of past phase-error signal, structurally identical to Definition 1’s capacitor voltage with V_d playing the role of current and the integrator output playing the role of v_C . The reason a Type-II PLL can do something a passive injection-locked oscillator cannot is now visible as a special case of this chapter’s general point: injection locking has no history-integrating state at all, while the Type-II loop filter gives the loop exactly the kind of accumulated-history state a capacitor has, and it is this added state that removes the bounded-locking-range limitation.*

Observation 38 (Feedback with reactive elements, deferred in Chapter 13). *Where a feedback path β includes capacitors or inductors, β is not a number but an operator built from the same history-integrals formalized in Definition 1, and the closed-loop output at any instant depends on the accumulated history the feedback path has integrated, not on the instantaneous input and output alone. Chapter 13’s fixed-point equation, $X_o = A(X_i - \beta X_o)$, is exact at every instant, but where β is history-dependent, the fixed point being solved for at time t is itself a function of everything that has passed through the reactive elements of the feedback path up to t .*

What This Chapter Does Not Claim

Observation 39. *Not every circuit element is history-defined, and this chapter's claim is specifically about capacitors, inductors, and the devices (like the diode's stored charge, or an active integrator) that share their integrative structure. A resistor's constitutive relation, $v = iR$, is purely algebraic: the current at an instant is a function of the voltage at that same instant, with no integral and no memory anywhere in the relation. An ideal op-amp's characteristic behavior (Chapter 11) is likewise memoryless in itself, apart from whatever reactive elements happen to surround it in a given circuit. Purely combinational digital logic (Chapter 17) is memoryless by design; it is sequential logic, built deliberately from memory elements, that reintroduces history into digital systems, and that distinction is itself the subject of a later chapter rather than a claim this one should anticipate. The claim of this chapter is not that circuits are history-defined everywhere; it is that a specific, identifiable, and practically central class of elements are, and that this book's vocabulary makes the shared structure across that class visible where standard treatment leaves each instance (capacitor, inductor, diode charge storage, loop-filter integrator) looking like an unrelated fact about a particular device.*

Observation 40. *It is also worth being precise about how simple a history a capacitor keeps, since not all of this book's historical objects are this simple. Capacitor charge accumulation is a linear, additive, non-branching compression: a single real number summarizing an integral, with no equivalent of an undo tree's preserved branches or a Vim macro's replayable structure. It is the most degenerate case of history-dependence available, which is precisely why it is the right place to start a chapter aiming to establish the claim rigorously rather than suggestively: if history-before-state did not hold even for the simplest possible historical object, it would be unlikely to hold for the richer ones.*

Verdict

This chapter's central result, Proposition 1, is a genuine proof rather than a restatement: it is not merely true that capacitor voltage "depends on" current history, it is provably the case that the map from history to state discards information, since distinct histories with equal accumulated charge are strictly indistinguishable afterward. This gives the book's recurring history-over-state theme its most rigorous grounding yet, in the most standard possible circuit element, and the observation that standard circuit theory already calls v_C and i_L "state variables" for exactly this reason means the chapter is making explicit an agreement that was already latent in ordinary engineering vocabulary rather than proposing a new one.

The chapter also discharges three specific debts rather than leaving them as unfulfilled promises: diode reverse recovery, the PLL integrator, and reactive feedback paths are all now shown to share one underlying structure, the same structure proven rigorously for the plain capacitor at the start of the chapter.

Histories Before States — Quick Reference

- * Capacitor/inductor constitutive relations, in integral form, are literally history-integrals: $v_C(t) = v_C(t_0) + \frac{1}{C} \int_{t_0}^t i \, d\tau$.
- * **Proposition:** distinct current histories with equal total charge yield identical $v_C(t_0)$ — the state provably discards information the history contained. Not a metaphor: a proof.
- * This is exactly why v_C, i_L are already called **state variables** in standard circuit/control theory — ordinary terminology already agrees with this book’s thesis.
- * **Three debts repaid:** diode reverse recovery ($Q_{rr} \approx \int i_F \, d\tau$, Ch. 9); PLL loop-filter integrator ($\int V_d \, d\tau$, PLL chapter); reactive feedback paths (Ch. 13) — all the same structure.
- * **Not universal:** resistors, ideal op-amps, and combinational logic are memoryless by design. History-dependence is a property of a specific, identifiable class of elements, not of circuits in general.
- * Capacitor history is the **simplest possible case** — linear, additive, non-branching — unlike undo trees or macros elsewhere in this author’s work. Deliberately the right place to start, not the richest example available.
- * **Verdict:** The book’s strongest grounding yet for its own central theme, proven on the simplest available element.

10 BJT/FET as Constrained Amplification Fields

Chapter 9 established the diode as this book’s first irreversible operator. A bipolar junction transistor’s base-emitter junction is, physically, the same kind of junction, so it is worth asking directly whether transistor amplification inherits Chapter 9’s irreversibility or belongs instead with the invertible operators of Chapters 5 and 6. The answer depends on which operating region is in play, and stating this precisely is this chapter’s main task.

10.1 The Active Region Is Invertible

Observation 41. *In the forward-active region, a BJT’s collector current follows $I_C = I_S e^{V_{BE}/V_T}$, the same Shockley-type exponential as the diode’s forward-bias equation. Restricted to this region, the function relating V_{BE} to I_C is smooth and strictly monotonic, hence invertible: given I_C , V_{BE} is uniquely recoverable. This is a materially different claim*

from Chapter 9's central result. Chapter 9's irreversibility concerned the idealized two-branch diode model, in which the entire reverse-bias half-line collapses to a single output; nothing in the active-region exponential collapses distinct inputs to a shared output. The active region belongs with the invertible operators of Chapters 5–6, not with Chapter 9's Pop/Refuse pair, despite both involving the same underlying exponential junction law.

Definition 18 (Small-signal transconductance). Linearizing $I_C = I_S e^{V_{BE}/V_T}$ about a bias point gives $i_c = g_m v_{be}$, with $g_m = \partial I_C / \partial V_{BE} \Big|_Q = I_C / V_T$, the small-signal transconductance.

Observation 42. At DC or low frequency, g_m is a real number: a pure scaling with no rotational component, a degenerate amplitwist operator with twist identically zero. Only at higher frequencies, where junction and parasitic capacitances become significant, does the small-signal transfer function acquire a genuine frequency-dependent phase, becoming a fuller amplitwist operator in the sense of Chapter 5. A low-frequency transistor small-signal model is therefore an amplitwist operator only in the most degenerate sense; the interesting case, where twist is actually nonzero, belongs to high-frequency transistor modeling this book does not pursue further.

10.2 Cutoff and Saturation: Where Collapse Returns

Observation 43. Outside the active region, the invertibility of Observation 1 fails, and it fails in a way that reproduces Chapter 9's structure rather than introducing a new one. In cutoff (V_{BE} below the conduction threshold), collector current is negligible regardless of exactly how far below threshold V_{BE} sits, a genuine Refuse-type collapse of an entire sub-threshold range to a single near-zero output. In saturation (both junctions forward-biased, V_{CE} small), collector current is no longer controlled by V_{BE} through the simple active-region relationship at all; the device output has, in effect, collapsed to a rail-like condition insensitive to further increases in base drive, a Pop-type collapse in Chapter 9's sense.

Corollary 3. A BJT is not uniformly one kind of operator. It is a device that is invertible (amplitwist-like) in its active region and Pop/Refuse-like (Chapter 9) at both of its boundaries, cutoff and saturation. Ordinary transistor amplifier design consists largely of ensuring the device is biased and driven so as to remain in the invertible region, deliberately avoiding the two regions where this book's irreversible vocabulary would otherwise apply.

10.3 The FET: A Different Nonlinearity, the Same Structure

Observation 44. A MOSFET in saturation follows a square law, $I_D \approx \frac{k}{2}(V_{GS} - V_{th})^2$, genuinely different from the BJT's exponential law; the two devices should not be treated as interchangeable nonlinearities differing only in packaging. The square law is likewise smooth and invertible on the domain $V_{GS} > V_{th}$, giving a small-signal transconductance $g_m = \partial I_D / \partial V_{GS} = k(V_{GS} - V_{th})$, again a real, degenerate amplitwist operator at low

frequency. An ideal MOSFET draws no DC gate current at all, unlike a BJT's nonzero base current, a distinguishing physical fact this book's Chapter 11 relies on directly.

10.4 Verdict

This chapter's contribution is a precise regional classification, not a claim that transistors are uniformly one kind of operator or another. The active (or saturation, for a FET) region is invertible and belongs with Chapters 5–6; cutoff and the BJT's saturation region are Pop/Refuse-like and belong with Chapter 9; ordinary linear amplifier design is, from this book's vocabulary, the discipline of keeping a device inside its invertible region and avoiding its two collapse regions.

11 Op-Amps as Idealized Admissibility Engines

An operational amplifier's ordinary analysis proceeds from two "golden rules": the two input terminals are at the same voltage (virtual short), and no current flows into either input. This chapter treats these rules not as approximations to be apologized for, but as a contextual admissibility criterion in the exact sense formalized in Chapter 4, and shows the criterion's own limit of exactness.

11.1 The Golden Rules as Contextual Admissibility

Observation 45. *An op-amp's open-loop behavior is $V_{out} = A_{ol}(V_+ - V_-)$, with A_{ol} typically 10^5 to 10^6 . If negative feedback holds V_{out} within the supply rails while $A_{ol} \rightarrow \infty$, then $(V_+ - V_-) = V_{out}/A_{ol} \rightarrow 0$ necessarily. The virtual-short golden rule is therefore not an independent postulate; it is the $A_{ol} \rightarrow \infty$ limit of an ordinary, finite-gain feedback amplifier, in exactly the fixed-point sense Chapter 13, later in this book, develops for feedback generally.*

Definition 19 (Ideal-op-amp admissibility criterion). *In the sense of Chapter 4, the ideal-op-amp analysis solves for the contextual admissibility of a circuit configuration relative to the criterion $C = \{V_+ = V_-, I_{in} = 0\}$.*

Observation 46. *For any real op-amp with finite A_{ol} , the admissibility gap $V_+ - V_- = V_{out}/A_{ol}$ is small but strictly nonzero. Ideal-op-amp analysis is exact only in the idealized limit; every specific circuit solved this way is being solved for the configuration admissible relative to C , with the finite-gain correction available, in principle, whenever A_{ol} is not large enough to be neglected.*

11.2 The Integrator: Where the Histories Chapter Becomes a Design Choice

Observation 47. *In an inverting integrator, virtual-ground admissibility forces the input resistor's current, V_{in}/R , to flow entirely into the feedback capacitor, giving $V_{out}(t) = -\frac{1}{RC} \int V_{in} d\tau$. This is the Histories-before-States chapter's capacitor constitutive relation, Definition 1 there, now harnessed deliberately as a computational primitive rather than encountered as an incidental physical fact about a passive element. The Histories chapter showed a capacitor accumulates history because it must; this circuit accumulates history because a designer wanted a circuit that integrates, and built one out of exactly the same physics.*

Observation 48. *The differentiator, $V_{out} = -RC \dot{V}_{in}$, is the dual operation: rather than accumulating history, it extracts the instantaneous rate of change, discarding everything about the signal except its present slope. Where the integrator is history-preserving in the strongest sense available in this book, the differentiator is close to the opposite: a memoryless functional of the input's local behavior, sensitive to noise for exactly this reason, since differentiation amplifies rapid, low-history fluctuations that integration would instead average away.*

11.3 The Comparator: Where Admissibility Gives Way to Collapse

Observation 49. *A comparator is an op-amp (or dedicated comparator device) operated open-loop, or with positive rather than negative feedback, so that $V_+ - V_-$ is not driven toward zero by the mechanism of Observation 5. Instead, an arbitrarily small input difference is amplified by the full open-loop gain until the output slams to one supply rail or the other. This is not a degraded or failed instance of the admissibility-engine regime of Sections 11.1–11.2; it is a different regime entirely, structurally identical to Chapter 9's Pop/Refuse pair and to Chapter 17's logic gates: an entire half-line of possible input differences (positive or negative) collapses to one of two discrete outputs. A comparator is, in this book's vocabulary, an analog-to-digital Pop operator, and clipping circuits built from op-amps and diodes inherit Chapter 9's diode vocabulary directly, limiting output excursion by the same forward/reverse collapse already established there.*

11.4 Verdict

The ideal-op-amp golden rules are shown here to be an admissibility criterion in the precise sense of Chapter 4, with a stated, quantifiable gap in the finite-gain case rather than an unexamined approximation. The integrator gives the Histories chapter's capacitor formalism its clearest deliberate engineering use in this book. The comparator shows that a device built from the same components as an integrator or a differentiator can belong to an entirely different regime, Chapter 9's collapse

rather than Chapter 4’s admissibility engine, depending only on whether negative feedback is present.

12 Log/Exponential Amplifiers as Compression Operators

A logarithmic amplifier, built from an op-amp with a diode or transistor junction in its feedback path, produces an output proportional to the logarithm of its input. This chapter asks what that logarithm is doing, in the same information-theoretic register this author’s other work uses for Assembly Index and Kolmogorov complexity, and finds a precise, quantitative answer rather than a loose resemblance.

12.1 The Log Amplifier as Dynamic-Range Compression

Observation 50. *Using the diode equation of Chapter 9 in an op-amp’s feedback path, virtual-ground admissibility (Chapter 11) gives $V_{out} = -\eta V_T \ln\left(\frac{V_{in}}{RI_S}\right)$. A tenfold change in V_{in} produces a fixed output change of $\eta V_T \ln 10 \approx 60$ mV per decade at room temperature for $\eta \approx 1$, independent of which decade is being crossed.*

Proposition 10 (Logarithmic compression of dynamic range). *Let V_{in} range over a dynamic range D (i.e. from some V_{min} to $D \cdot V_{min}$). The corresponding output range is $\eta V_T \ln D$, growing with the logarithm of D rather than with D itself.*

Proof. Immediate from Observation 7: the output at V_{min} and at $D \cdot V_{min}$ differ by $\eta V_T \ln D$, regardless of the absolute scale of V_{min} . \square

Observation 51. *This is a precise, quantitative instance of the same principle underlying this author’s Assembly-Theoretic treatment of Kolmogorov complexity elsewhere, where representing a choice among N equally likely possibilities costs $\log_2 N$ bits rather than N resources. A logarithmic amplifier is a physical, continuous-domain circuit realizing exactly this compression: an output range requirement scaling with $\ln D$ rather than D is the analog-circuit counterpart of a bit cost scaling with $\log_2 N$ rather than N . This is a structural and quantitative correspondence, not merely two unrelated uses of the same function; both are instances of logarithms being the natural cost of representing a large space of distinguishable magnitudes or possibilities in a bounded output range. It is not a claim that a log amplifier computes an Assembly Index or a Kolmogorov complexity in any operational sense; it realizes the same compression principle in a different, physical domain.*

12.2 The Log-Antilog Multiplier and Chapter 6’s Composition Law

Observation 52. *An analog multiplier can be built from the log amplifier of this chapter by a route worth deriving correctly: two log converters produce $\eta V_T \ln(V_1/RI_S)$ and*

$\eta V_T \ln(V_2/R I_S)$ separately; summing these (an ordinary op-amp summing junction, Chapter 11) gives an output proportional to $\ln(V_1/R I_S) + \ln(V_2/R I_S) = \ln\left(\frac{V_1 V_2}{(R I_S)^2}\right)$; a final exponential (antilog) amplifier stage then recovers a signal proportional to the linear product $V_1 V_2$. This is the standard log-antilog (translinear) analog multiplier technique, and it is worth being precise that this, not a single-stage circuit producing $\ln(V_1) \ln(V_2)$ directly, is the mathematically correct route: summing two exponentials does not algebraically reduce to a product of their arguments' logarithms, and no single summing-junction stage computes $\ln(V_1) \ln(V_2)$ from V_1 and V_2 directly.

Observation 53. *The log-antilog multiplier's principle, converting multiplication into addition via a logarithm and back again via an exponential, is the identical algebraic move underlying Chapter 6's matched-cascade composition law, where transmission-coefficient magnitudes multiply under cascading and this is why engineers add decibels (logarithms of those magnitudes) rather than multiplying raw ratios. Chapter 6 used the log-linearization of amplitude composition as a purely mathematical convenience for reasoning about cascaded stages; this chapter's log-antilog multiplier is the same convenience realized as an actual analog computing circuit, converting a multiplication a designer wants performed into two additions and two logarithmic nonlinearities a diode or transistor junction already provides for free.*

12.3 Where This Chapter's Idealizations Fail

Observation 54. *Three limitations are worth naming. First, $V_T = k_B T/q$ is temperature-dependent, so the log and antilog scale factors drift with temperature unless deliberately compensated, a standard practical concern in real log-amplifier design. Second, the ideal exponential diode law of Chapter 9 holds only over a finite current range in a real device; series resistance and other non-idealities cause deviation from the pure logarithmic relationship at both very low and very high input levels, bounding the dynamic range over which Proposition 1 holds exactly. Third, the Assembly-Theoretic correspondence of Section 12.1 is a structural and quantitative analogy about compression in general, not a claim that any specific Assembly Index or Kolmogorov complexity value is being computed by the circuit; the connection should be read at the level of the shared principle, not as an operational equivalence.*

12.4 Verdict

This chapter's central results are a precise proposition (logarithmic compression of dynamic range, Proposition 1) and a corrected engineering derivation (the log-antilog multiplier, replacing an erroneous single-stage shortcut with the standard, correct translinear technique). The connection to this author's Assembly-Theoretic vocabulary is stated at the level it can actually support, a shared compression principle rather than an operational identity, and the connection to Chapter 6's composi-

tion law is exact rather than suggestive: both are the same log-converts-multiplication-to-addition move, one used as mathematical bookkeeping, the other built out of silicon.

13 Feedback as Recursive Continuation

The synchronization and phase-locked loop chapters both assume, by name, a treatment of feedback and of the Barkhausen criterion that this chapter and the next now supply directly, since both were written to rely on grounding they had not yet been given. This chapter takes the more basic case first: what feedback is, formally, before any question of oscillation arises.

13.1 The Closed Loop as a Fixed-Point Equation

Definition 20 (Feedback loop). *Let a forward path have gain A and a feedback path have gain β , with input X_i and output X_o . A negative-feedback loop defines X_o implicitly by*

$$X_o = A(X_i - \beta X_o).$$

Observation 55. *This equation does not compute X_o from X_i in one forward pass; it defines X_o as whatever value satisfies its own defining equation, i.e. as a fixed point of the map $g(X_o) = A(X_i - \beta X_o)$. This is not a metaphorical use of “recursive continuation”; it is the literal mathematical structure of the closed loop. Solving,*

$$X_o = \frac{A}{1 + A\beta} X_i,$$

gives the familiar closed-loop gain, but the familiar formula is the resolved form of a genuinely self-referential equation, not evidence that the self-reference was dispensable.

This is, among the correspondences drawn in this book so far, one of the least effortful: no population, no order parameter, no approximation regime is required to see a feedback amplifier as an instance of a system defined by the requirement that it continuously regenerate the condition of its own output. The algebra makes the point without needing to be pushed toward it.

13.2 Feedback in Time: Self-Reference as an ODE

Observation 56. *In the time domain, a feedback loop with reactive elements is a differential equation in which the rate of change of the output depends on the output itself, mediated by the feedback path: schematically $X_o = h(X_o, X_i)$ for some h depending on X_o through β . The static fixed-point picture of Section 13.1 is the equilibrium of this dynamic picture; stability of that equilibrium is a separate question, taken up in Chapter 14, and is not guaranteed merely because a fixed point exists algebraically. A fixed point that is unstable is a real solution of the static equation that the physical system will never settle into.*

13.3 A Brief, Deferred Note on History

Observation 57. *Where the feedback path includes reactive elements, β is not merely a number but an operator on the signal's past (a filter), and the loop's present output then depends on accumulated history rather than on the present input and output alone. This is the same fact this book's Histories-before-States chapter treats in full for capacitors, inductors, and the phase-locked loop's own integrating filter; it is noted here only as a pointer, since a full treatment belongs to that chapter and would be premature here.*

14 The Barkhausen Criterion as an Admissibility Condition

Chapter 13 established feedback as a fixed-point equation, with stability of that fixed point left open. This chapter treats the specific boundary at which the fixed point stops being a stable equilibrium and a self-sustained oscillation becomes possible instead, which is the condition the synchronization and phase-locked loop chapters both invoked by name in advance of this treatment.

14.1 The Criterion, Recalled

Definition 21 (Barkhausen criterion). *A feedback loop with open-loop gain $A(j\omega)\beta(j\omega)$ satisfies the Barkhausen criterion at frequency ω_0 if*

$$|A(j\omega_0)\beta(j\omega_0)| = 1 \quad \text{and} \quad \angle A(j\omega_0)\beta(j\omega_0) = 0 \pmod{360^\circ}.$$

14.2 The Criterion Is the Linear Onset Condition for a Hopf Bifurcation

Proposition 11. *The closed-loop characteristic equation of the feedback loop in Definition 1 (Chapter 13) is $1 - A(s)\beta(s) = 0$. The Barkhausen criterion at real frequency ω_0 is exactly the condition that this characteristic equation has a root at $s = j\omega_0$, i.e. a pair of complex-conjugate closed-loop poles located precisely on the imaginary axis.*

Proof. $1 - A(j\omega_0)\beta(j\omega_0) = 0$ if and only if $A(j\omega_0)\beta(j\omega_0) = 1$, which holds if and only if both $|A(j\omega_0)\beta(j\omega_0)| = 1$ and its phase is a multiple of 360° , which is Definition 1. □

Observation 58. *A pair of poles crossing the imaginary axis as a parameter is varied is precisely the linear-stability signature of a Hopf bifurcation in nonlinear dynamics: the point at which a stable equilibrium loses stability and a limit cycle becomes possible. This is not an analogy imported from this author's prior work on locomotion; the Barkhausen*

critera, developed in circuit theory independently of bifurcation theory, and the Hopf onset condition, developed in dynamical systems theory independently of circuit engineering, are describing the identical linear-algebraic event, a pair of eigenvalues (poles) crossing $j\omega$. The correspondence noted in this author's locomotion volume between CPG onset and Hopf bifurcation and the correspondence noted here between Barkhausen onset and Hopf bifurcation are two independent instances of the same mathematical fact, not two applications of one borrowed metaphor.

14.3 Necessary, Not Sufficient: Startup Requires More

Observation 59. *The Barkhausen criterion as stated in Definition 1 describes a marginal condition: poles sitting exactly on the imaginary axis correspond to a linear system that neither grows nor decays, and a real physical system starting from rest (or from noise) at exactly this condition has no reason to depart from zero amplitude, since the condition is one of neutral stability, not instability. A self-starting oscillator requires $|A(j\omega_0)\beta(j\omega_0)| > 1$ at small signal, placing the closed-loop poles strictly in the right half-plane so that noise or a turn-on transient grows; amplitude-dependent gain compression (saturation, limiting, or any other gain-reducing nonlinearity) then reduces the effective loop gain as amplitude grows, until the effective $|A\beta|$ falls back to exactly 1 at some finite, self-limited amplitude. The Barkhausen criterion, taken literally and linearly, describes the steady-state amplitude condition, not the startup condition, and stating it as sufficient for oscillation, as introductory treatments frequently do, elides this distinction.*

Corollary 4. *This is the precise circuit-theoretic content behind this book's earlier use, in the oscillator-circuits material presupposed by the synchronization chapter, of Hopf language for oscillation onset: small-signal instability ($|A\beta| > 1$, poles in the right half-plane) together with amplitude-dependent self-limiting nonlinearity is the supercritical Hopf picture (soft, self-starting growth from zero amplitude) rather than the marginal, literally-stated Barkhausen condition alone, which corresponds only to the bifurcation point itself, not to the mechanism that gets a real system across it.*

14.4 Admissibility, Named Precisely

Corollary 5. *The admissibility framing invoked by name in the synchronization chapter (sustained oscillation as an admissible, continuously maintained condition rather than a discrete achieved state) and in the phase-locked loop chapter (lock as continuously regenerated coherence) rests on exactly the mechanism made precise in this section: oscillation persists exactly as long as the small-signal instability and the self-limiting nonlinearity continue to jointly satisfy the effective Barkhausen condition at the operating amplitude, moment to moment, not as a condition established once and then assumed to hold.*

14.5 Where This Chapter's Picture Breaks Down

Observation 60. *Three limitations are worth naming. First, a real amplifier's phase response can satisfy the phase condition of Definition 1 at more than one frequency; when this happens, whichever frequency reaches $|A\beta| = 1$ first, or most robustly, tends to dominate, but multi-mode competition and mode-hopping are real, documented complications the single-frequency treatment here does not capture. Second, component tolerance and drift mean the exact phase condition is never satisfied with perfect precision in a physical circuit, so real oscillators operate with margin rather than at an exact boundary, a practical point this idealized treatment elides. Third, and most fundamentally, the linear Barkhausen analysis of this section is blind to the steady-state oscillation amplitude entirely; amplitude is set exclusively by the nonlinear self-limiting mechanism referenced in Section 14.3, and determining it requires nonlinear methods (such as the describing-function method) outside the scope of the linear criterion this chapter has formalized.*

14.6 Verdict

The central claim of this pair of chapters is not that feedback and oscillation onset can be described in this book's vocabulary, but that two of the results obtained are independent, checkable facts rather than restatements. Chapter 13's fixed-point reading of feedback requires no approximation and no population; it is the literal content of the closed-loop equation. This chapter's identification of the Barkhausen criterion with the linear onset condition for a Hopf bifurcation is a genuine correspondence between two independently developed theories (classical feedback circuit analysis and nonlinear dynamical systems theory), not a metaphor built for this book. The necessary-but-not-sufficient distinction of Section 14.3 is a real correction to a common oversimplification, in the same register as Chapter 5's addition-versus-composition distinction, and the admissibility framing used by the two later chapters that presupposed this one is now grounded rather than merely asserted.

Chapters 13–14 Quick Reference

- * **Feedback** ($X_o = A(X_i - \beta X_o)$) is a literal fixed-point equation: X_o is defined by requiring it to satisfy its own defining condition, resolved as $X_o = \frac{A}{1+A\beta} X_i$. This is exact, not metaphorical, recursive continuation.
- * **Barkhausen criterion:** $|A\beta| = 1$, phase = $0^\circ \pmod{360^\circ} \iff$ closed-loop poles exactly on the $j\omega$ axis — the identical linear event as a **Hopf bifurcation** onset in dynamical systems theory. Independently developed, genuinely convergent.
- * **Necessary, not sufficient:** literal Barkhausen is a marginal (steady-state) condition. Self-starting requires $|A\beta| > 1$ at small signal (poles in RHP) plus amplitude-dependent gain compression bringing the loop back to $|A\beta| = 1$ at a finite, self-limited amplitude — the supercritical-Hopf picture, precisely.
- * This grounds the **admissibility** language already used in the Synchronization and PLL chapters: oscillation persists only as long as this condition is jointly, continuously satisfied.
- * **Breaks down:** multi-frequency phase matches (mode competition); component tolerance (no exact boundary in practice); amplitude is entirely outside this linear treatment's scope (needs nonlinear/describing-function analysis).
- * **Verdict:** Two independent, checkable results (fixed-point feedback; Barkhausen=Hopf onset), not restatements — and a real correction to a common textbook oversimplification.

15 Oscillator Circuits as Self-Sustaining Amplitwist Loops

Chapters 13 and 14 developed the general theory this chapter now instantiates: feedback as a fixed-point equation, and the Barkhausen criterion as the linear onset condition for a Hopf-type bifurcation, with startup requiring small-signal instability and steady-state amplitude set by self-limiting nonlinearity. Chapters 13–14 referred to “the oscillator-circuits chapter” by name in advance of this one; this chapter now supplies the concrete circuits that reference presupposed, and is written to be consistent with, rather than a revision of, what was already assumed about it. The Synchronization chapter, later in this book, in turn refers back to this chapter's material when it takes up multiple, coupled oscillators.

15.1 The General Structure, Recalled and Specialized

Proposition 12 (Two conditions, evaluated together). *For any oscillator built from a feedback loop with frequency-dependent open-loop gain $A(j\omega)\beta(j\omega)$, the oscillation frequency ω_0 is determined by the phase condition alone, $\angle A(j\omega_0)\beta(j\omega_0) = 0 \pmod{360^\circ}$, and the minimum gain required to sustain oscillation is determined by evaluating the magnitude condition, $|A(j\omega_0)\beta(j\omega_0)| \geq 1$, at that same frequency ω_0 , not at any other frequency.*

This is simply Definition 1 of Chapter 14 applied in the order an actual design proceeds: solve for the frequency at which the network's phase shift is right, then ask how much gain is needed at that particular frequency. Each oscillator topology below is an instance of this same two-step structure, differing only in what feedback network $\beta(j\omega)$ supplies the phase shift.

15.2 The RC Phase-Shift Oscillator

A three-section RC ladder feedback network, combined with an inverting amplifier stage, is among the simplest oscillators in this sense: the amplifier alone supplies 180° of the required phase shift, and the RC network must supply the remaining 180° at the frequency of oscillation.

Observation 61. *Because each RC section loads the section before it (the sections are not buffered from one another), the three-section network's phase and magnitude response cannot be obtained by simply cubing a single section's transfer function; the loaded ladder must be analyzed as a single third-order network. This is a standard, well-established derivation (see, e.g., Sedra and Smith, or Boylestad, for the full nodal analysis), and this chapter reports its result rather than re-deriving the loaded-ladder algebra from first principles here, since the intermediate steps are lengthy and the result is not in question: for three identical sections with resistance R and capacitance C , the phase condition is satisfied at*

$$\omega_0 = \frac{1}{RC\sqrt{6}}, \quad f_0 = \frac{1}{2\pi RC\sqrt{6}},$$

and evaluating the magnitude condition at this same ω_0 shows the loaded network attenuates the signal by a factor of 29 at this frequency, so the amplifier stage must supply a gain magnitude of at least 29 to satisfy Proposition 1's magnitude condition.

Observation 62. *The point worth extracting from this example, independent of the specific numbers $\sqrt{6}$ and 29, is the logical structure Proposition 1 makes explicit: the frequency and the minimum gain are not two independent design choices, they are two properties of the same feedback network evaluated at the single frequency the phase condition selects. A designer who chose an amplifier gain based on the network's attenuation at some other frequency, or who tuned R and C without checking that the resulting ω_0 was the frequency*

at which the required gain was evaluated, would be solving the two Barkhausen conditions independently rather than jointly, and would generally fail to obtain a working oscillator even with algebraically correct component values.

15.3 The LC (Hartley) Oscillator

The Hartley oscillator supplies its phase shift through an LC resonant tank rather than an RC ladder, and here the derivation is short enough to give in full.

Proposition 13 (Hartley resonance condition). *For a tank formed by two inductors L_1, L_2 with mutual inductance M and a capacitor C , the tank's reactance vanishes, and the transistor stage's own phase shift alone is sufficient to satisfy the Barkhausen phase condition, at*

$$\omega_0 = \frac{1}{\sqrt{C(L_1 + L_2 + 2M)}}, \quad f_0 = \frac{1}{2\pi\sqrt{C(L_1 + L_2 + 2M)}}.$$

Proof. The effective series inductance of the tapped inductor pair, including their mutual coupling, is $L_{\text{eff}} = L_1 + L_2 + 2M$; setting the standard LC resonance condition $\omega_0 L_{\text{eff}} = \frac{1}{\omega_0 C}$ and solving for ω_0 gives the stated result directly. \square

Observation 63. *This is a direct instance of the point this book's amplitwist reading of resonance makes: at ω_0 , the tank's impedance operator degenerates from a general rotation-and-scaling to a pure scaling (its reactive, rotation-contributing part vanishing exactly), which is precisely why the tank contributes no additional phase shift at this frequency and the Barkhausen phase condition reduces to whatever phase the active device alone contributes. Resonance, in an LC oscillator, is not merely a frequency-selective filter effect; it is the specific condition under which the feedback network's own amplitwist operator stops rotating, leaving the loop's phase budget to be satisfied by the amplifying device alone.*

15.4 The Crystal Oscillator

Observation 64. *A crystal oscillator (Pierce configuration or similar) replaces the LC tank with a piezoelectric resonator, modeled, near its series resonance, by an equivalent motional inductance L_s and capacitance C_s , giving an oscillation frequency of approximately $f_0 \approx \frac{1}{2\pi\sqrt{L_s C_s}}$, structurally the same LC resonance condition as Proposition 2, with the crystal's mechanical resonance standing in for a discrete inductor and capacitor.*

Observation 65. *The crystal's practical advantage, exceptional frequency stability, is directly explained by machinery the Synchronization chapter, immediately following this one, develops in full; it is previewed briefly here rather than derived. That chapter's locking-range formula, $|\Delta\omega| \leq \frac{\omega_0}{2Q}\varepsilon$, shows locking range shrinking as Q grows; a crystal's very high Q (commonly 10^4 to 10^6 , orders of magnitude above a typical LC tank) means that whatever spurious coupling or injected interference a real circuit is exposed to, the resulting locking*

range within which that interference could pull the crystal's frequency is correspondingly narrow. A crystal oscillator resists frequency pulling not because of some property unrelated to injection locking, but because its high Q places it deep in the regime where that formula, given in full in the next chapter, predicts resistance to being locked by anything except a very close, very strong interfering signal.

15.5 Amplitude Self-Limiting, Concretely

Observation 66. Chapter 14 established, in the abstract, that self-starting oscillation requires $|A\beta| > 1$ at small signal and an amplitude-dependent nonlinearity to bring the effective loop gain back to exactly 1 at some finite amplitude. In each circuit above, the concrete mechanism is ordinary transistor (or op-amp) gain compression: as oscillation amplitude grows, the active device begins to approach its supply rails or its linear operating region's edge, its incremental gain falls, and the effective $|A\beta|$ decreases until it reaches unity at the amplitude where growth stops. No separate limiting circuit is required in the simplest designs; the same nonlinearity that would ordinarily be considered distortion is exactly the mechanism Chapter 14 required to exist.

15.6 Where This Chapter's Idealizations Break Down

Observation 67. Three limitations are specific to these concrete circuits rather than to the abstract theory of Chapters 13–14. First, the RC phase-shift result of Section 15.2 assumes a non-loading ideal amplifier input and output; a real amplifier's finite input and output impedance perturbs the loaded-ladder analysis and shifts both the frequency and the minimum required gain from the idealized values reported there. Second, component tolerance and temperature drift affect R , C , L directly, and therefore ω_0 directly, in the RC and LC cases; the crystal oscillator's advantage is precisely that its resonance depends on a stable mechanical property rather than on ordinary passive component tolerances, which is why crystal references, not LC tanks, are used wherever long-term frequency accuracy matters. Third, gain compression sufficient to self-limit amplitude, relied on in Section 15.5, can in some designs be abrupt enough (hard clipping rather than gradual compression) to introduce harmonic distortion or, in poorly designed loops, more complex nonlinear behavior beyond simple amplitude limiting; the clean supercritical-Hopf picture of gradual self-limiting is a good first approximation, not a guarantee for every practical implementation.

15.7 Verdict

This chapter does not introduce new abstract machinery; its task was to show that the Barkhausen/Hopf theory of Chapters 13–14, already used by name in two later chapters before it was itself written, correctly organizes three standard, independently well-established oscillator circuits under a single two-condition structure. The RC phase-shift result is reported rather than re-derived, honestly, to avoid the

risk of a memory error in lengthy loaded-ladder algebra; the LC (Hartley) result is derived in full, since it is short and safe; and the crystal oscillator's practical advantage is explained, not merely asserted, by direct appeal to the locking-range formula the Synchronization chapter develops immediately afterward. The chapter closes the specific gap Chapters 13–14 had been presupposing, without needing to revise anything they had already claimed.

Synchronization: From Coupled Oscillators to Network Coherence

The oscillator-circuits chapter preceding this one treats a single feedback loop crossing the Barkhausen condition, $|A\beta| = 1$ with zero total phase shift, as the birth of a sustained oscillation: a single ampliwist operator, in the Hopf sense clarified in this author's prior work on locomotion, acquiring a nonzero limit-cycle amplitude as loop gain crosses the threshold. That chapter concerns exactly one oscillator. This chapter concerns what happens when two or more such oscillators are permitted to influence one another, and exists specifically to prepare the ground for the Phase-Locked Loop chapter that follows it, in the same way the biological and Kuramoto background of this author's locomotion volume was placed before, rather than folded into, the CPG chapters that depended on it.

From One Oscillator to Two: Injection Locking

The relevant phenomenon, well established in electronics independent of anything in this book, is injection locking: a free-running oscillator, when a sufficiently strong external signal at a nearby frequency is injected into it, abandons its own natural frequency and locks to the injected frequency instead.

Definition 22 (Injection-locked phase equation, after Adler). *Let an oscillator with natural (free-running) angular frequency ω_0 and quality factor Q be injected with an external signal at frequency ω_{inj} , with injection ratio ε (the ratio of injected signal amplitude to the oscillator's own steady-state amplitude, assumed small). Let $\phi(t)$ be the phase difference between the oscillator and the injected signal, and $\Delta\omega = \omega_0 - \omega_{inj}$ the detuning. Adler's equation governs the evolution of ϕ :*

$$\dot{\phi} = \Delta\omega - \frac{\omega_0}{2Q} \varepsilon \sin \phi.$$

Observation 68. *This is not offered as an analogy to the Kuramoto phase equation used throughout this author's prior work; it is the identical equation. Writing $K = \frac{\omega_0}{2Q} \varepsilon$, Adler's equation reads $\dot{\phi} = \Delta\omega - K \sin \phi$, which is exactly the two-oscillator Kuramoto equation in the frame co-rotating with the injected reference. This correspondence is documented in the*

synchronization literature independent of this book (Adler’s original 1946 result predates Kuramoto’s 1975 model by three decades, and the equivalence between the two is standard background in later treatments of both). Circuit engineers derived, and have used for eighty years, a special case of the equation this book has built two chapters around, without needing the word “amplitwist” to do so.

Proposition 14 (Locking range). *An injection-locked oscillator governed by Adler’s equation reaches a stable fixed point (phase lock, $\dot{\phi} = 0$) if and only if*

$$|\Delta\omega| \leq \frac{\omega_0}{2Q}\varepsilon,$$

the locking range. Outside this range, no fixed point exists and $\phi(t)$ drifts monotonically (with nonuniform angular velocity), corresponding to unlocked beating between the oscillator and the injected signal rather than synchronization.

Proof. Setting $\dot{\phi} = 0$ in Adler’s equation requires $\sin \phi = \frac{2Q\Delta\omega}{\omega_0\varepsilon}$, which has a solution in $[-1, 1]$ exactly when $|\Delta\omega| \leq \frac{\omega_0}{2Q}\varepsilon$. When a solution exists, the fixed point with $\cos \phi > 0$ is stable (standard linearization of Adler’s equation about the fixed point); when no solution exists, $\dot{\phi}$ never vanishes and ϕ increases or decreases without bound, modulo 2π . \square

This is the direct circuit analogue of the Kuramoto critical coupling K_c that governs the onset of synchronization in a population of oscillators: the locking range is a two-oscillator critical coupling, derived from ordinary linear circuit quantities (Q , injection ratio) rather than posited abstractly.

From Two Oscillators to a Population

The same phase-coupling structure extends, with the same order parameter used throughout this author’s work on locomotion, to arrays of many mutually coupled oscillators, a genuine and actively used technique in microwave engineering for power combining and beam steering (coupled oscillator arrays, following York and Compton).

Definition 23 (Coupled oscillator array). *Let $\{\phi_i(t)\}_{i=1}^N$ be the phases of N nominally identical oscillators, each coupled to some set of neighbors with coupling strength K_{ij} :*

$$\dot{\phi}_i = \Delta\omega_i + \sum_j K_{ij} \sin(\phi_j - \phi_i).$$

Define the array’s order parameter exactly as in the Kuramoto reduction used throughout this author’s prior work,

$$Z(t) = r(t)e^{i\psi(t)} = \frac{1}{N} \sum_{i=1}^N e^{i\phi_i(t)}.$$

Observation 69. *Nothing in this definition is new relative to this author's treatment of locomotor CPGs; the equation is the same equation, and $Z(t)$ is the same order parameter, with ϕ_i now the phase of an electrical oscillator rather than a limb's phase oscillator. The substantive content of this chapter is not a new formalism; it is the observation that this formalism was already independently required by electrical engineers building synchronized oscillator arrays, for reasons having nothing to do with locomotion or with this book.*

Synchronization as an Admissibility Condition

The Barkhausen chapter preceding this one already frames sustained oscillation as an admissibility condition: a loop either satisfies $|A\beta| = 1$ with zero phase shift, and oscillation continues, or it does not, and oscillation dies out or grows without bound. Locking extends this same admissibility framing from a single oscillator's continuation to a population's joint continuation.

Corollary 6. *An injection-locked oscillator, or a coupled oscillator array, is in an admissible synchronized state exactly when the coherence measure of Definition 2 (single-pair: bounded phase difference per Proposition 1; population: $r(t)$ bounded away from a critical value) is maintained. Loss of lock, whether of a single injection-locked oscillator falling outside its locking range or of a coupled array's $r(t)$ collapsing, is a loss of admissibility in precisely the sense developed for balance in this author's work on locomotion: not a discrete failure event alone, but the crossing of a continuously monitored coherence threshold.*

Bridge to the Phase-Locked Loop

The Phase-Locked Loop chapter that follows this one is, in light of the preceding sections, best introduced as an engineered generalization of injection locking rather than as an unrelated circuit topology.

Observation 70. *An injection-locked oscillator is coupled to its reference passively, through whatever parasitic or intentional coupling path carries the injected signal, with a locking range fixed by the oscillator's own Q and the available injection strength. A phase-locked loop replaces this passive coupling with an engineered, actively filtered feedback path: a phase detector explicitly computes a phase-error signal, a loop filter shapes it, and a voltage-controlled oscillator's own free-running frequency is steered by the resulting control voltage rather than left fixed. The PLL is what injection locking becomes once the coupling term is no longer a fixed physical parameter but a designed transfer function the engineer can shape. Its locking range and lock dynamics, treated in the next chapter, are governed by equations of the same phase-coupled family as Adler's equation and the Kuramoto array of this chapter, now with the coupling strength itself made a function of frequency through the loop filter's design.*

Where This Chapter's Equations Break Down

Observation 71. *Adler's equation is a weak-injection, narrowband approximation, and this chapter's clean correspondence to the Kuramoto equation inherits exactly that restriction. For strong injection ratios, the oscillator's amplitude can no longer be treated as approximately fixed, higher-order locking to subharmonics and superharmonics of the injected frequency becomes possible, and the simple locking-range formula of Proposition 1 is replaced by a considerably more intricate structure of locking regions in the (detuning, injection strength) plane, known as Arnold tongues, with chaotic dynamics possible between them (documented in the large-signal injection-locking literature, e.g. Kurokawa). Separately, the array treatment of Section 15.2 assumes coupling close to all-to-all or otherwise well-mixed; real coupled-oscillator arrays are frequently coupled only to physical neighbors, in which case the mean-field order parameter of Definition 2 is a less exact description, and spatially structured states (phase gradients, traveling phase waves across the array) can appear in place of uniform lock. Both limitations mirror, in a different physical setting, the mean-field-versus-microscopic caveat already stated for the locomotion volume's Kuramoto reduction.*

Verdict

This chapter's central claim is not a new derivation but a documented equivalence: Adler's equation, independently derived by electrical engineers in 1946 for entirely practical reasons, is the same equation this author's prior work derives from coupled biological oscillators, and the coupled-oscillator-array literature in microwave engineering already uses the identical order parameter this book has used throughout. This is, if anything, a stronger form of evidence than either Chapter 5's forced conceptual distinction or Chapter 6's derived composition law: it is not this book finding a place where its formalism applies, but a demonstration that two independent engineering and biological literatures converged on the same equation without reference to each other, which this book is positioned to notice rather than to have caused. The honest limitation — that this correspondence holds in a weak-coupling, mean-field regime, and both assumptions fail in identifiable, named ways — keeps the chapter from overstating what has been shown.

16 The Phase-Locked Loop as Engineered Continuation

The preceding chapter closed by describing the phase-locked loop as injection locking with its passive coupling replaced by an engineered, filtered feedback path. This chapter makes that claim precise, shows it holds exactly in the simplest case, and then asks the question that gives the chapter its reason to exist: what does

engineering the coupling actually buy, beyond what a passively injection-locked oscillator already had? If the answer were nothing, this chapter would be a re-labeling of the last one. It is not nothing, and showing why is this chapter's central task.

16.1 The Loop, in Block Form

A phase-locked loop consists of a phase detector (PD) producing an error signal from the phase difference between a reference input and the loop's own output, a loop filter (LF) shaping that error signal into a control voltage, and a voltage-controlled oscillator (VCO) whose instantaneous frequency is steered by that control voltage. Writing θ_{ref} and θ_{vco} for the reference and VCO phases, and $\phi = \theta_{\text{ref}} - \theta_{\text{vco}}$ for the phase error:

$$\begin{aligned} V_d &= K_d \sin \phi && \text{(phase detector)} \\ V_c &= F(s) V_d && \text{(loop filter)} \\ \dot{\theta}_{\text{vco}} &= \omega_{\text{free}} + K_{\text{vco}} V_c && \text{(VCO)} \end{aligned}$$

16.2 The Simplest Case Is Exactly Adler's Equation

Proposition 15 (Type-I PLL recovers the injection-locking equation). *For a Type-I loop, in which the loop filter is a pure gain, $F(s) = F_0$, the phase-error dynamics reduce to*

$$\dot{\phi} = \Delta\omega - K_d K_{\text{vco}} F_0 \sin \phi, \quad \Delta\omega = \omega_{\text{ref}} - \omega_{\text{free}},$$

which is Adler's equation exactly, with the effective coupling $K = K_d K_{\text{vco}} F_0$ now a product of engineered gains rather than a fixed physical injection ratio.

Proof. Differentiating $\phi = \theta_{\text{ref}} - \theta_{\text{vco}}$ and substituting the VCO equation with $V_c = F_0 K_d \sin \phi$ gives $\dot{\phi} = \dot{\theta}_{\text{ref}} - \omega_{\text{free}} - K_{\text{vco}} F_0 K_d \sin \phi = \Delta\omega - K \sin \phi$ directly. \square

Corollary 7. *A Type-I PLL has exactly the locking range of Proposition 1 in the previous chapter, $|\Delta\omega| \leq K = K_d K_{\text{vco}} F_0$, now expressed in terms of chosen circuit gains rather than an oscillator's Q and an injection ratio set by the physical coupling path. Nothing about the underlying phase dynamics has changed; only the source of the coupling strength has.*

This is the sense in which the bridge asserted at the end of the previous chapter holds exactly rather than loosely: a Type-I PLL is not merely similar to an injection-locked oscillator, it is one, with an engineered rather than parasitic coupling constant.

16.3 What Engineering the Coupling Buys: An Integrator

The genuine advance a PLL offers over passive injection locking is available only once the loop filter is given dynamics of its own, and the clearest case is a Type-II loop filter with an integrating term, $F(s) = \frac{1+s\tau_2}{s\tau_1}$.

Observation 72. *Introducing an integrator into the loop filter changes the system's order. Adler's equation, and the Type-I PLL of Proposition 1, are first-order phase equations: a single state variable ϕ fully determines the dynamics. A Type-II loop filter carries its own internal state (the integrator's accumulated output), so the loop as a whole is a second-order system: phase error and the filter's internal state jointly determine the dynamics. This is not a minor refinement; it is the addition of a genuine second degree of freedom the injection-locked oscillator of the previous chapter did not have.*

Proposition 16 (Zero steady-state phase error under a frequency step). *For a Type-II PLL, linearized for small phase error ($\sin \phi \approx \phi$) about lock, the open-loop transfer function contains two integrators: one from the VCO (phase is the integral of frequency) and one from the loop filter itself. By the final value theorem, the steady-state phase error in response to a step change in reference frequency (a ramp in reference phase) is*

$$\phi(\infty) = \lim_{s \rightarrow 0} s \cdot \frac{\Delta\omega/s^2}{1 + K_d K_{vco} F(s)/s} = 0,$$

since the $1/s$ term contributed by the loop filter's integrator forces the denominator's low-frequency behavior to cancel the $1/s^2$ of the ramp input exactly, whereas a Type-I loop (a single open-loop integrator, from the VCO alone) leaves a nonzero steady-state error for the same input, and loses lock altogether once $|\Delta\omega|$ exceeds the fixed range of Corollary 1.

Observation 73. *This is the chapter's central result, and it is worth stating plainly what kind of advance it is. An injection-locked oscillator, however strongly driven, has a locking range fixed by its physical parameters; detuning beyond that range is unrecoverable without changing the hardware. A Type-II PLL, within its VCO's tuning range, can track an arbitrarily large constant frequency offset with zero steady-state phase error, because the integrator continuously adjusts the control voltage until the VCO's actual frequency matches the reference exactly, however large the initial offset. Engineering the coupling did not just relocate the constant K of Adler's equation into a design parameter; adding a second, historical degree of freedom removed the bounded-lock-range limitation altogether, for this class of input.*

This is also, without needing to force the connection, an instance of this book's recurring theme that circuit elements are frequently defined by accumulated history rather than instantaneous state: the loop filter's integrator output at any moment is the time integral of all past phase-error signals, not a function of the present phase error alone, and it is precisely this historical accumulation that lets the loop achieve what a memoryless (Type-I, first-order) coupling cannot.

16.4 A Second Extension: Capture Range

Observation 74. *A separate limitation of the sinusoidal phase detector, $V_d = K_d \sin \phi$, is that its usable range for pulling an unlocked loop into lock (the capture or pull-in range) is generally narrower than its locking range once already locked, because far from lock the averaged beat-note error signal the detector produces is small and slowly varying. Practical PLLs frequently use a phase-frequency detector (PFD) instead, which produces an output sensitive to the sign of the frequency difference as well as the phase difference, dramatically widening the capture range relative to a simple multiplying or XOR-based phase detector. This is a large-signal, strongly nonlinear phenomenon, well documented in the standard PLL design literature (e.g. Best), and is not derivable from the small-phase-error linearization used in Proposition 2; it is noted here as a real and practically important effect this chapter's linear analysis does not, and does not claim to, capture.*

16.5 Failure Modes Specific to This Chapter

Observation 75. *Three failure modes are worth naming, distinct from the Arnold-tongue and mean-field limitations of the previous chapter, because they arise specifically from the added dynamics introduced in this one. First, cycle slipping: if the phase error is driven beyond $\pm\pi$ (typically during acquisition from a large initial frequency offset, before lock is achieved), the loop effectively loses an entire reference cycle before re-attempting lock, a discrete, topological event in phase space rather than a smooth departure from a fixed point, and one the continuous linearization of Proposition 2 does not describe. Second, loop stability: a Type-II loop's damping ratio ζ , determined by the loop filter's time constants together with $K_d K_{vco}$, can be driven low enough by aggressive gain or filter choices to produce ringing or outright instability in the closed-loop response, an ordinary control-theoretic risk introduced precisely by the second-order dynamics that made Proposition 2 possible in the first place. Third, false or harmonic lock: a real phase detector's periodic characteristic can permit the loop to settle at a stable phase relationship corresponding to a harmonic or subharmonic of the intended reference frequency rather than the fundamental, a failure mode with an evident family resemblance to the Arnold-tongue subharmonic locking noted for strongly driven injection-locked oscillators in the previous chapter, now arising from the phase detector's own periodicity rather than from strong injection.*

16.6 Verdict

Does this chapter show something beyond the previous one? Yes, and the demonstration is a genuine derived result rather than a relabeling: Proposition 1 shows the Type-I PLL is exactly the injection-locking equation of the previous chapter, with an engineered rather than physical coupling constant, establishing the bridge asserted there was not merely apt but exact. Proposition 2 then shows what engineering the coupling can add once the loop filter is allowed internal dynamics: a

strictly stronger tracking property (zero steady-state error, unbounded-in-principle frequency tracking within the VCO's range) that a bare injection-locked oscillator, however strongly coupled, structurally cannot achieve, because it has no analogue of the integrator's accumulated history.

Does the chapter overclaim? It should not. The zero-steady-state-error result is specific to a frequency step and to the linearized, small-phase-error regime; the capture-range extension from a PFD is noted honestly as a separate, large-signal phenomenon this chapter's linear analysis cannot derive; and cycle slipping, loop instability, and harmonic lock are named as concrete limits on when either result actually holds in a physical loop.

17 Digital Logic as Spherpocalypse

Chapter 9 established the diode as this book's first genuinely irreversible operator, and proved that rectification specifically requires non-invertibility. It also declined, deliberately, to invoke Landauer's bound for that device, since a diode is a continuous, unlocked component and the bound concerns discrete, clocked logical erasure. This chapter is where that declined connection becomes available in earnest: a combinational logic gate is exactly the kind of discrete, well-defined, typically clocked operation Landauer's argument addresses, and this chapter both proves the necessity of collapse for standard gates and quantifies it.

17.1 Every Standard Logic Gate Is a Pop Operator

Proposition 17. *Every Boolean function $f : \{0, 1\}^2 \rightarrow \{0, 1\}$ is non-injective.*

Proof. The domain $\{0, 1\}^2$ has four elements; the codomain $\{0, 1\}$ has two. By the pigeonhole principle, at least two distinct inputs must share an output value. \square

Corollary 8. *Every two-input, one-output logic gate—AND, OR, NAND, NOR, XOR, XNOR, without exception—is a Pop operator in the sense of Chapter 9: a bounded region (the pair of input values) collapsing to a value (the single output bit) that does not, by itself, determine which input produced it.*

This holds for every such gate regardless of which specific function it computes, which raises the natural next question: do different gates destroy the same amount of information, or different amounts?

17.2 Not All Collapses Are Equal

Definition 24 (Information destroyed by a gate). *For a gate $Y = f(X)$ with X uniformly distributed over its input space, the information destroyed by the gate is the conditional entropy $H(X | Y)$: the residual uncertainty about which input occurred, given only the output.*

Proposition 18. For AND (equivalently OR, NAND, NOR, by symmetry of the calculation), $H(X | Y) \approx 1.189$ bits. For XOR (equivalently XNOR), $H(X | Y) = 1$ bit exactly.

Proof. For X uniform over $\{0, 1\}^2$, $H(X) = 2$ bits. For AND, $P(Y = 0) = 3/4$ (satisfied by three of the four inputs) and $P(Y = 1) = 1/4$ (satisfied only by $(1, 1)$), giving $H(Y) = -\frac{3}{4} \log_2 \frac{3}{4} - \frac{1}{4} \log_2 \frac{1}{4} \approx 0.811$ bits. Since Y is a deterministic function of X , $H(X) = H(Y) + H(X | Y)$, so $H(X | Y) = 2 - 0.811 \approx 1.189$ bits. For XOR, $P(Y = 0) = P(Y = 1) = 1/2$ (each satisfied by exactly two inputs), giving $H(Y) = 1$ bit exactly and $H(X | Y) = 2 - 1 = 1$ bit. \square

Observation 76. Every two-input gate is non-invertible, by Proposition 1, but Proposition 2 shows the gates are not equally destructive of information: AND destroys more than XOR does, for uniformly random inputs. This is a genuine, quantitative refinement of the Pop-operator picture, not merely a restatement that “gates lose information.” It also connects this chapter directly to the entropy and distinguishability material of Chapters 20 and 21: information destroyed by a gate is, in the same units, the same kind of quantity as the distinguishability threatened by noise in those chapters, though the mechanism here is exact logical collapse rather than continuous fluctuation.

17.3 Landauer’s Bound, Properly Earned Here

Observation 77. Landauer’s principle states that erasing one bit of information in a physical system carries an unavoidable minimum thermodynamic cost of $k_B T \ln 2$, dissipated as heat, regardless of how the erasure is physically implemented. Chapter 9 declined to invoke this bound for the diode, since the bound concerns discrete, clocked logical erasure and a diode is a continuous, unclocked device. A logic gate operating within a synchronous digital system is precisely the setting the bound was developed for: a well-defined discrete input state is replaced, at a definite clocked instant, by a well-defined discrete output state that does not determine the input. This chapter is entitled to the connection Chapter 9 declined.

Proposition 19 (Minimum dissipation per gate operation). A gate destroying $H(X | Y)$ bits of information per operation (Definition 1, Proposition 2) has a minimum thermodynamic dissipation, per operation, of $k_B T \ln 2 \cdot H(X | Y)$.

Observation 78. This gives AND a minimum dissipation of approximately $1.189 k_B T \ln 2$ per operation, and XOR a minimum of exactly $k_B T \ln 2$, under uniform random inputs. This is a genuine, quantitative lower bound, not an estimate of actual gate dissipation. Real CMOS logic gates dissipate several orders of magnitude more energy per operation than this fundamental limit, for reasons having entirely to do with practical transistor switching losses rather than with logical necessity; Proposition 3 states a floor, not a prediction of measured power consumption.

17.4 Reversible Gates: The Exception That Proves the Rule

Definition 25 (Reversible logic gate). *A logic gate is reversible if its Boolean function is a bijection on its full input space, so that every output determines its input uniquely.*

Observation 79. *Standard two-input, one-output gates cannot be reversible, by Proposition 1; a reversible gate necessarily has equal numbers of input and output bits (e.g. the three-input, three-output Toffoli gate, or the Fredkin gate). Bennett showed that computation can, in principle, be carried out entirely with reversible gates, and that such computation carries no Landauer-bound dissipation at all, since $H(X | Y) = 0$ for any bijection. This gives the digital domain a clean extension of the reversible/irreversible split already running through this book: reversible gates belong with the invertible amplitude operators of Chapters 5 and 6 (zero forced dissipation), and standard AND/OR/NAND/NOR-type gates belong with the Pop operators of Chapter 9 (a strictly positive dissipation floor, quantified by Proposition 3).*

17.5 Sequential Logic: Where History Re-enters

Observation 80. *Combinational logic, the subject of this chapter so far, is memoryless by design: each output is a function of the present inputs alone, exactly as the Histories-before-States chapter noted in passing without developing further. Sequential logic elements, flip-flops and latches, reintroduce history deliberately: a D flip-flop's output Q after a clock edge depends on the value of D at the previous clock edge, not on any combinational input at the present instant. Q is, in the precise sense developed for capacitor voltage in the Histories chapter, a state: a compressed summary of a trajectory (the sequence of past clocked inputs) that discards everything about that sequence except its most recent clocked value. The discrete case is, if anything, a cleaner illustration of the same principle than the continuous capacitor was: a flip-flop's state is a single bit, and the entire history prior to the most recent clock edge is provably irrelevant to future behavior, not merely compressed but discarded outright.*

17.6 Failure Modes Specific to This Chapter

Observation 81. *Three limitations are worth naming. First, Proposition 3's dissipation floor assumes uniformly random inputs; real digital workloads are typically correlated and non-uniform, so the information actually destroyed, and the corresponding Landauer floor, is workload-dependent rather than a fixed property of the gate alone. Second, and more interestingly, a flip-flop is physically implemented as a bistable feedback circuit, typically a pair of cross-coupled gates forming a positive feedback loop with two stable equilibria corresponding to its two logic states; when input timing constraints (setup and hold time) are violated, the circuit can be driven toward the unstable equilibrium between these two stable states and linger there, in a metastable condition, for an unpredictable duration before resolving unpredictably to one state or the other. This is precisely the instability structure*

introduced abstractly in Chapter 13: a flip-flop's normal operation relies on two stable fixed points of a feedback equation, and metastability is the practical failure mode of landing near the unstable fixed point that necessarily separates them. Third, real logic gates have nonzero propagation delay and are not perfectly synchronized to an idealized clock edge, so the clean discrete-time picture of Section 17.5 is, like every idealization in this book, a limit rather than an exact description of physical timing.

17.7 Verdict

This chapter's central results are genuine proofs and a genuine quantitative refinement, not restatements. Proposition 1 establishes, with a two-line pigeonhole argument, that every standard logic gate is non-invertible without exception. Proposition 2 shows this non-invertibility is not uniform across gates, giving AND and XOR provably different information-theoretic costs. Proposition 3 connects that quantitative difference to an actual physical dissipation floor, in a setting (discrete, clocked digital logic) where Chapter 9 explicitly declined to make the analogous claim for a continuous device. The reversible-gate material shows the irreversible/reversible split already used for analog operators in Chapters 5, 6, and 9 extends cleanly into the digital domain, and the sequential-logic material cashes a promise deferred explicitly by the Histories-before-States chapter. The metastability failure mode is not an afterthought; it is a direct, concrete instance of Chapter 13's abstract feedback-stability machinery, arising naturally once flip-flops are examined as the physical feedback circuits they are rather than as idealized black boxes.

18 Clock Distribution, Skew, and Domain Crossing

The Synchronization chapter treated coherence among coupled oscillators as a statistical, phase-coupling phenomenon, governed by an order parameter and a critical coupling strength. Clock distribution within a single digital system is, for the most part, a different kind of problem entirely: a deterministic propagation-delay problem, not a statistical one, and this chapter is deliberately scoped to avoid re-deriving the Synchronization chapter's machinery where it does not apply, while using it explicitly, by citation, where it does.

18.1 Clock Skew Is Deterministic, Not Statistical

Definition 26 (Clock skew). *Let a single clock signal be distributed to two points in a circuit, arriving at times t_1 and t_2 for corresponding edges. The skew between these points is $\delta = t_2 - t_1$, a quantity determined by the physical propagation delay of the distribution network (wire length, buffer stages, capacitive loading) to each point.*

Observation 82. *Skew is not a coherence deficit in the Kuramoto sense; it is not a statistical property of a population of oscillators drifting in and out of phase alignment, but a fixed (to first order) consequence of physical layout. A clock tree with badly matched path lengths has large skew regardless of how perfectly stable and noise-free its source oscillator is; this is a different failure mode from anything the order parameter of the Synchronization chapter describes, and this chapter’s early sections should not be read as a special case of that chapter’s material.*

18.2 Setup and Hold as a Joint Admissibility Criterion

Correct operation of a register-to-register path in a synchronous system requires two timing inequalities to hold simultaneously, one guarding against data arriving too late (a setup constraint) and one guarding against data arriving too early (a hold constraint). The exact derivation of both, including sign conventions for skew, is standard and available in any digital design text (e.g. Harris and Harris); this chapter states their general form and focuses its own contribution on what they jointly mean in this book’s admissibility vocabulary.

Observation 83. *In general form, a setup constraint takes the shape $T_{\text{clock}} \geq t_{CQ} + t_{\text{logic,max}} + t_{\text{setup}} - \delta$, and a hold constraint takes the shape $t_{CQ,\text{min}} + t_{\text{logic,min}} \geq t_{\text{hold}} + \delta$, for skew δ defined as in Definition 1. The two constraints respond to skew in opposite directions: skew that relaxes the setup constraint (by giving the destination more effective time before its capturing edge) simultaneously tightens the hold constraint (by giving old data less effective margin before being overwritten), and vice versa. Skew cannot be made to help both constraints on the same path at once.*

Proposition 20 (Skew-induced admissibility can be empty). *For a fixed clock period T_{clock} and fixed logic delays, there exist values of skew δ for which no assignment of the remaining free design parameters satisfies both the setup and hold constraints of Observation 2 simultaneously.*

Proof. The setup constraint requires $\delta \leq t_{CQ} + t_{\text{logic,max}} + t_{\text{setup}} - T_{\text{clock}}$ and the hold constraint requires $\delta \leq t_{CQ,\text{min}} + t_{\text{logic,min}} - t_{\text{hold}}$ (rearranging Observation 2’s two inequalities to isolate δ with consistent sign). If the physical layout forces δ above both of these bounds simultaneously, for instance through an unavoidably long clock-tree path difference, no choice of T_{clock} within otherwise-acceptable operating limits resolves both inequalities at once, and the admissible set of Chapter 4 (Definition 2 there) is empty for this path. \square

Observation 84. *This is a second, independent instance of Chapter 4’s admissibility-emptiness result, now in a digital-timing domain rather than the oscillator-design domain that result was originally stated for. Timing closure failure in real chip design, where a design is abandoned or re-laid-out because no clock period and no skew budget can be found that satisfies every register-to-register path simultaneously, is exactly this proposition occurring in practice, not a metaphorical use of the word “admissible.”*

18.3 Clock Domain Crossing: Where Chapter 17's Metastability Is Unavoidable

Definition 27 (Clock domain crossing). *A signal crosses clock domains when it passes from logic clocked by one clock to logic clocked by a second clock with no fixed, known phase relationship to the first.*

Observation 85. *Sections 18.1–18.2 concern paths where skew, though it may be large, is at least a fixed, known, computable quantity, so that the two timing inequalities of Observation 2 can in principle be checked. A clock domain crossing admits no such computation: since the two clocks have no fixed phase relationship, the destination flip-flop's capturing edge can fall anywhere relative to the source signal's transition, on any given crossing. Chapter 17 established that a flip-flop is physically a bistable feedback circuit that can be driven toward its unstable equilibrium when setup or hold timing is violated; a clock domain crossing is precisely the case in which no timing analysis can guarantee this never happens, because there is no fixed timing relationship left to analyze. Every individual crossing carries some nonzero probability of metastability, not as a design flaw to be eliminated, but as an unavoidable consequence of the two clocks' independence.*

Observation 86. *The standard mitigation is a multi-flop (commonly two-flop) synchronizer: the crossing signal is captured by a first flip-flop in the destination domain, which may momentarily go metastable, but is then given a full additional clock cycle to resolve before a second flip-flop samples its output. The probability that metastability persists longer than the available resolution time falls off exponentially with the time allowed, governed by the destination flip-flop technology's own metastability resolution time constant; this is the basis of standard mean-time-between-failures calculations used in real synchronizer design (see Chaney and Molnar for the original characterization of this behavior). A two-flop synchronizer does not eliminate the possibility of a metastable value propagating; it reduces the probability to a level considered acceptable for the application, which is a different and weaker guarantee.*

18.4 Distributed Synchronization: A Direct Application of an Earlier Chapter

Observation 87. *Not every clock distribution problem is a tree-and-skew problem. Some systems, particularly those spanning multiple boards, chassis, or independently clocked subsystems, distribute timing by mutual synchronization among multiple oscillators rather than by buffering a single source through a hierarchical tree, using techniques (GPS-disciplined oscillators, mutually locked PLL networks) that are direct, uncomplicated applications of the Synchronization chapter's Kuramoto and Adler-equation machinery, already fully developed there. This chapter does not re-derive that material; where distributed timing genuinely is a coherence problem among independently running oscillators, the correct reference is the Synchronization chapter itself, not a restatement of it here.*

18.5 Failure Modes and Limitations

Observation 88. *Three further limitations are worth naming. First, this chapter’s treatment of skew as static and deterministic omits jitter, cycle-to-cycle timing variation in the clock’s own edges, which is a genuinely statistical quantity, physically driven in large part by the same thermal and shot-noise mechanisms formalized in Chapter 20, now manifesting as timing uncertainty rather than amplitude uncertainty; a complete treatment of clock timing margins must budget for jitter in addition to the static skew this chapter has focused on. Second, the exponential metastability-resolution argument of Section 18.3 describes an ever-decreasing probability, not a guarantee, and safety-critical designs must choose a synchronizer resolution time appropriate to the failure rate the application can tolerate, not assume metastability has been eliminated outright. Third, transferring more than one bit across a clock domain boundary introduces a distinct failure mode beyond single-bit metastability: different bits of a multi-bit value can be sampled correctly as individual bits while still representing a torn, inconsistent combination if the value changes at the wrong moment relative to the destination clock, a problem simple per-bit synchronizers do not address and which instead motivates techniques such as Gray coding or handshake-based transfer protocols.*

18.6 Verdict

This chapter deliberately avoided re-deriving the Synchronization chapter’s coherence machinery, since clock skew and setup/hold timing are a different, deterministic problem, correctly distinguished in Section 18.1. Its genuine technical contribution is Proposition 1, a second and independent instance of Chapter 4’s admissibility-emptiness result, now shown to hold in ordinary digital timing closure rather than only in the oscillator-design context that result was first stated for. Clock domain crossing gave Chapter 17’s metastability material its first genuinely unavoidable application, rather than a rare edge case, and Section 18.4 discharged, by explicit citation rather than duplication, the one place where this chapter’s subject matter and the Synchronization chapter’s actually do coincide.

19 Memory and Repair

The Histories-before-States chapter established that a capacitor’s voltage is a compressed, lossy summary of its current history. Chapter 17 showed that a flip-flop’s stored bit is a discrete instance of the same idea. Neither chapter asked what happens when that stored information is threatened, actively, by decay or corruption, and what it takes to keep it admissible over time. This chapter treats digital memory specifically as a place where this author’s prior, more general treatment of repair, developed for telemetry and long-timescale concept formation, has a precise and checkable circuit-level instance.

19.1 DRAM as Continuously Repaired History

A dynamic random-access memory cell stores a bit as charge on a small capacitor, exactly the device formalized in the Histories chapter, but with one addition that chapter did not need: leakage.

Definition 28 (Leaky memory cell). *A DRAM cell's stored voltage decays according to*

$$V(t) = V_0 e^{-(t-t_{\text{last}})/\tau}, \quad \tau = R_{\text{leak}}C,$$

where t_{last} is the time of the most recent write or refresh, and R_{leak} models the cell's unavoidable non-infinite leakage resistance.

Observation 89. *This is the Histories chapter's capacitor with an additional dissipative path, and the dissipative path is not incidental to this chapter's concerns: leakage is exactly the kind of entropy-producing process formalized in Chapter 20, now acting directly against the persistence of stored information rather than merely producing heat as a byproduct. A DRAM cell is where the Histories chapter's "state as compressed history" and Chapter 20's "dissipation as entropy production" meet directly: the cell's bit is a history under continuous thermodynamic erosion.*

Proposition 21 (Maximum safe refresh interval). *Let V_{th} be the minimum voltage at which a stored 1 remains reliably distinguishable from a stored 0. The cell remains admissible (correctly readable) only for $t - t_{\text{last}} < t_{\text{max}} = \tau \ln(V_0/V_{\text{th}})$. A refresh policy with interval t_{refresh} maintains continuous admissibility if and only if $t_{\text{refresh}} < t_{\text{max}}$.*

Proof. Setting $V(t) = V_{\text{th}}$ in Definition 1 and solving for $t - t_{\text{last}}$ gives t_{max} directly. If refresh (a read followed by a rewrite to V_0 , resetting t_{last}) occurs at intervals shorter than t_{max} , the voltage never reaches V_{th} before being reset. \square

Observation 90. *This is a literal, physical instance of recursive continuation as developed elsewhere in this author's work: a DRAM bit persists not because it is stored once and left alone, but because an active process (refresh) continuously regenerates the condition (sufficient charge) for its own future readability. Real DRAM refresh intervals are commonly specified around 64 ms (JEDEC-standard values vary by device and temperature); the entire discipline of memory refresh circuitry exists because Proposition 1's inequality would otherwise be violated by ordinary leakage on ordinary timescales.*

19.2 Error-Correcting Memory as the Repair Chain, Concretely

This author's prior work on procedural ontology formalizes repair as a composition, $R(\Gamma) = \text{Sort}_{\prec_R} \circ \text{Validate}_{\sim_V}(\Gamma)$: a validation relation \sim_V first determines which fragments are structurally consistent, and a distinct ordering or correction relation \prec_R then reconstructs an admissible object from what validation has passed. Error-correcting memory is a clean, worked instance of exactly this two-step chain, not a loose analogy to it.

Definition 29 (Hamming(7,4) code, recalled). *Four data bits are encoded into seven transmitted bits by adding three parity bits, each computed over a specific, overlapping subset of positions determined by the binary representation of the position indices 1 through 7. On receipt, three parity checks are recomputed; their results, read as a binary number, form the syndrome $s \in \{0, \dots, 7\}$.*

Proposition 22. *If $s = 0$, the received word is a valid codeword (no single-bit error). If $s \neq 0$, s gives the position of the single erroneous bit, which repair corrects by flipping it.*

Observation 91. *This is the abstract repair chain instantiated exactly. Syndrome computation is the validation relation \sim_V : it determines, from the received fragment Γ (the possibly-corrupted seven-bit word), whether and where a structural inconsistency exists, without yet fixing anything. Bit-flipping at the indicated position is the distinct repair operation proper: it reconstructs an admissible codeword from what validation identified. Chapter 25 of this author's procedural ontology work insisted on keeping these two operations conceptually separate rather than treating repair as a single undifferentiated act; Hamming decoding is a place where the separation is not merely conceptually clean but is literally how the algorithm is implemented, syndrome computation and bit correction being distinct steps in every practical decoder.*

19.3 Silent Miscorrection: When Repair Produces a Wrong, Admissible History

Observation 92. *Hamming(7,4) corrects any single-bit error, but if two bits are corrupted, the syndrome computed by Proposition 2 is generally nonzero and points to some position, real or spurious, and the decoder will "correct" that position with full confidence. The result is a codeword that is admissible, in the narrow sense of satisfying all parity checks, but is not the codeword that was originally transmitted. The decoder reports success; the repair has, in fact, produced a coherent-looking but incorrect history. This is not a design flaw specific to Hamming(7,4); it is a necessary consequence of a code with a fixed minimum distance being asked to correct more errors than that distance guarantees.*

Observation 93. *This failure mode deserves to be named precisely because it recurs, in different clothing, throughout this author's other work on repair: admissibility and correctness are not the same property, and a repair operation can produce an object that passes every validation check available to it while still being the wrong object. Practical memory systems address this specific failure by adding redundancy beyond single-error-correction, most commonly an additional overall parity bit giving SECDED (single-error-correction, double-error-detection) codes, which cannot correct a double error but can at least flag that one has occurred rather than silently miscorrecting it. The lesson generalizes beyond circuits: a repair mechanism is only as trustworthy as the gap between the errors it is designed to correct and the errors it is actually asked to face.*

19.4 Failure Modes and Limitations

Observation 94. *Three limitations are worth stating plainly. First, Proposition 1’s refresh guarantee assumes a leakage time constant τ that is itself vulnerable to environmental variation; elevated temperature reduces R_{leak} and hence τ , and radiation-induced charge disturbances can cause abrupt, non-leakage-related bit upsets that no refresh schedule timed for ordinary leakage will anticipate. Second, Hamming-style codes are designed against independent, randomly located bit errors; real memory faults are frequently correlated or burst in nature (an entire row or column of a memory array failing together), against which a code designed for isolated single-bit errors offers little protection, motivating different code families (such as Reed–Solomon codes) explicitly designed for burst-error environments. Third, and most generally, Section 19.3’s miscorrection failure mode is not a rare edge case to be dismissed; it is the generic behavior of any fixed-distance code once the actual error rate exceeds its design assumptions, and treating a passed validation check as a guarantee of correctness, rather than merely an absence of detected inconsistency, is the specific epistemic error this section warns against.*

19.5 Verdict

This chapter’s central contribution is not a new formalism but a demonstration that a formalism already developed elsewhere in this author’s work, the repair chain $R(\Gamma) = \text{Sort}_{\sim_R} \circ \text{Validate}_{\sim_V}(\Gamma)$, has an exact, checkable circuit-level instance in Hamming decoding, with syndrome computation and bit correction mapping onto validation and repair respectively without needing to be forced into correspondence. The DRAM refresh material gives recursive continuation its most literal circuit-level instance in this book: a bit persists only because an active process continuously regenerates the physical condition for its own future readability, against a genuine, quantified thermodynamic decay. The miscorrection failure mode is this chapter’s most important honest limitation, not because it is unique to memory systems, but because it names precisely the gap, familiar from this author’s broader repair-theoretic work, between an object passing validation and an object being correct.

20 Noise as Entropic Production

Chapter 9 noted, in passing, that forward diode conduction is unavoidably dissipative, and that this was the first point in this book at which the entropy field S of the RSVP triad (Φ, v, S) became unavoidable rather than a background assumption. That chapter declined to make a specific quantitative claim about entropy at the time, reserving it for here. This chapter makes good on that reservation, and does so by anchoring S to a real, quantitative, independently established physical

result, the fluctuation-dissipation theorem, rather than to a formal analogy built only on this book's own vocabulary.

20.1 Entropy Production in a Dissipative Circuit

Definition 30 (Entropy production rate of a resistor). *For a resistor with resistance R carrying current I at absolute temperature T , dissipating power $P = I^2 R$ as heat into its surroundings, the rate of thermodynamic entropy production is*

$$\frac{dS}{dt} = \frac{P}{T} = \frac{I^2 R}{T} \geq 0.$$

Proposition 23 (A Second Law for circuits). *For any passive network composed of positive resistances R_i carrying currents I_i , together with ideal (lossless) reactive elements, the total rate of entropy production is*

$$\frac{dS}{dt} = \frac{1}{T} \sum_i I_i^2 R_i \geq 0,$$

with equality if and only if no current flows through any resistive element.

Proof. Each term $I_i^2 R_i \geq 0$ by positivity of R_i and of I_i^2 ; the ideal reactive elements contribute no dissipative term by definition (Chapters 5–6 and the Histories chapter treat them as lossless in the ideal case), so the sum is nonnegative, with equality exactly when every $I_i = 0$ through every resistive element. \square

Observation 95. *This is the literal, physical content underlying the “Second Law for Code” proven in this author’s prior work on procedural ontology, where entropy production for an executing program was established formally within a Shannon-entropy, execution-path framework. Proposition 1 is not an application of that earlier, more abstract result; it is a separate, independent instance of the same inequality, this time in actual thermodynamic entropy, actual temperature, and actual dissipated power, requiring no appeal to this book’s own formalism to be true. That two structurally similar Second Laws hold, one for program execution and one for dissipative circuits, is worth noting without overstating: both are consequences of entropy production being nonnegative wherever genuine dissipation occurs, which is a general thermodynamic fact, not a coincidence specific to either domain.*

20.2 The Same Resistors That Dissipate Also Fluctuate

The chapter’s central result connects Proposition 1 to something apparently unrelated: the noise voltage measurable across an ordinary resistor sitting at thermal equilibrium, with no current source driving it at all.

Theorem 2 (Johnson–Nyquist noise). *A resistor of resistance R at absolute temperature T exhibits a mean-square thermal noise voltage, in a measurement bandwidth Δf , of*

$$\overline{v_n^2} = 4k_B T R \Delta f,$$

independent of any current deliberately driven through the resistor.

Observation 96. *This result, established independently by Johnson and Nyquist in 1928, is the fluctuation-dissipation theorem applied to a resistor: the same parameter R that determines how much power a resistor dissipates when current is driven through it (Definition 1) also determines how strongly that resistor’s own thermal agitation fluctuates when no external current is applied at all. Dissipation and fluctuation are not two independent facts about a resistor that happen to share a symbol; the fluctuation-dissipation theorem shows they cannot be varied independently; a hypothetical resistor with the dissipative behavior of Definition 1 but without the thermal noise of Theorem 1 would violate this theorem and could, in principle, be exploited to violate the second law of thermodynamics. Chapter 9’s dissipative diode and this chapter’s noisy resistor are, at the deepest level available in this book, the same physical fact viewed from two directions: entropy production is dissipation, and dissipation is fluctuation.*

20.3 Shot Noise: A Second, Distinct Mechanism

Observation 97. *A second noise mechanism, distinct from Johnson–Nyquist thermal noise, is relevant specifically to the diode of Chapter 9 and to other junction devices: shot noise, arising from the discreteness of charge carriers crossing a potential barrier, with mean-square current fluctuation $\overline{i_n^2} = 2qI \Delta f$ (after Schottky’s original 1918 treatment), where q is the electron charge and I the average current. Unlike Johnson–Nyquist noise, shot noise does not vanish as $T \rightarrow 0$; it is a consequence of charge discreteness, not of thermal agitation, and its physical origin is genuinely different from the fluctuation-dissipation mechanism of Theorem 1, even though both are commonly grouped together as “noise” in circuit design practice. This chapter’s entropy-production framing applies cleanly to Johnson–Nyquist noise, where dissipation and fluctuation are theorem-linked; it does not extend automatically to shot noise, which this chapter is careful not to claim.*

20.4 Reactive Elements Are, Ideally, Silent

Observation 98. *The elements identified in this book’s Histories-before-States chapter as accumulating history without dissipation, ideal capacitors and inductors, are correspondingly noiseless in the same ideal limit: an ideal, lossless reactive element has no resistive term to appear in Proposition 1’s sum, and correspondingly no thermal noise source in Theorem 1’s sense. This is a clean, consistent split running through several chapters of this book at once: the elements that are history-defined and reversible (Chapters 5–6, and the Histories chapter) are, ideally, also the elements that neither dissipate nor fluctuate, while*

the elements that are irreversible in the Pop/Refuse sense of Chapter 9 are exactly the elements this chapter identifies as unavoidably both dissipative and noisy. Real capacitors and inductors have parasitic series resistance and are, to that extent, noisy in practice; the ideal case is a limit, not a claim about physical components, in the same way Chapters 5 and 6 treated ideal losslessness as a limit throughout.

20.5 Where This Chapter's Picture Is Incomplete

Observation 99. *Two limitations are worth naming plainly. First, flicker noise ($1/f$ noise), ubiquitous in real active devices at low frequencies, lacks as clean a first-principles derivation as either Johnson–Nyquist or shot noise; its physical origins are heterogeneous and only partially understood in general, and this chapter does not attempt an entropy-production account of it. Second, the fluctuation-dissipation theorem invoked in Theorem 1 is, in its simplest form, an equilibrium or near-equilibrium result; devices operating far from equilibrium (transistors under large-signal drive, for instance) require a more general nonequilibrium statistical mechanics to justify a fluctuation-dissipation-style relationship rigorously, and this chapter's clean equilibrium treatment should not be assumed to extend there without further argument.*

20.6 Verdict

This chapter's central claim rests on real, independently established physics, not on an extension of this book's own vocabulary: the fluctuation-dissipation theorem, and its specific circuit instance in Johnson–Nyquist noise, is a genuine theorem relating dissipation and fluctuation, proven in statistical mechanics without reference to anything in this book. What this chapter contributes is the observation that this theorem gives the entropy field S , used loosely elsewhere in this book's vocabulary, an exact and quantitative circuit-theoretic referent for the first time: entropy production is literally dissipated power over temperature, and the same resistive elements responsible for it are, by an independent theorem, the same elements responsible for thermal noise. The chapter is honest about where this clean picture stops: shot noise is a distinct mechanism not reducible to the same theorem, flicker noise resists a comparably clean first-principles account altogether, and the equilibrium assumption underlying Theorem 1 is not automatically available far from equilibrium.

21 Signal-to-Noise Ratio and the Compactness Criterion

Chapter 20 established that dissipation and fluctuation are linked by the fluctuation-dissipation theorem. A resistor that dissipates necessarily fluctuates; noise is not

an engineering inconvenience layered atop a fundamentally noiseless substrate but a direct consequence of the same physical processes that permit dissipation in the first place. That chapter deliberately stopped at the existence of noise. This chapter asks a different question: how much useful structure survives once noise is present?

Ordinary engineering answers this question through signal-to-noise ratio. This chapter argues that signal-to-noise ratio can be understood as a quantitative version of a more general idea already present throughout this book: the distinction between admissible structure and background variability.

21.1 Signal and Noise as Competing Components

Let a measured signal be decomposed as $x(t) = s(t) + n(t)$, where $s(t)$ is the intended signal and $n(t)$ is noise. The standard engineering definition of signal-to-noise ratio is

$$\text{SNR} = \frac{P_s}{P_n}, \quad \text{SNR}_{dB} = 10 \log_{10} \left(\frac{P_s}{P_n} \right),$$

where P_s and P_n are the average signal and noise powers respectively.

Observation 100. *Nothing in these definitions is unique to this book; they are among the most familiar quantities in communication theory and electronic measurement. The question is not whether SNR exists, but what structural role it plays.*

21.2 The Recoverability Interpretation

Definition 31 (Recoverable signal). *A signal component is recoverable if it can be distinguished from noise with reliability exceeding a specified threshold.*

Observation 101. *Signal-to-noise ratio is therefore not fundamentally a measure of signal strength; it is a measure of distinguishability. A weak signal with extremely low noise may be easier to recover than a strong signal buried beneath larger fluctuations. The quantity being measured is not energy, but distinguishability.*

21.3 Distinguishability and Admissibility

Chapter 8 interpreted filtering as an admissibility mechanism. Signal-to-noise ratio introduces a second admissibility condition.

Definition 32 (Detection admissibility). *A signal is detection-admissible if its signal-to-noise ratio exceeds the minimum threshold required for reliable recovery by a specified detector.*

Observation 102. *The threshold is not universal; different receivers, algorithms, and tasks require different values. Admissibility remains contextual rather than absolute, exactly as it did in Chapter 8. A filter asks whether a frequency is admissible; a detector asks whether a distinction is admissible. These are different questions with the same logical structure.*

21.4 Noise as Compression Pressure

Observation 103. *Noise acts as compression pressure: as noise increases, more distinct signal configurations become observationally indistinguishable from one another. Two signals that differ only slightly remain distinguishable at low noise and become observationally equivalent at high noise, at which point a receiver cannot determine which was originally present.*

Proposition 24. *Increasing noise reduces the number of reliably distinguishable signal states available to a fixed receiver.*

Proof. As noise variance increases, the probability distributions associated with different signal states overlap more strongly. Increased overlap decreases classification reliability and increases ambiguity between states. \square

Observation 104. *This is one reason communication engineers devote so much effort to noise reduction: noise does not merely corrupt values, it collapses distinctions.*

21.5 Shannon Capacity as an Admissibility Volume

Theorem 3 (Shannon capacity). *For a channel of bandwidth B and signal-to-noise ratio SNR,*

$$C = B \log_2(1 + \text{SNR})$$

is the maximum achievable information rate for arbitrarily reliable communication.

Observation 105. *This theorem is not derived here; it is stated as an anchor connecting the admissibility language of earlier chapters to a standard, independently established engineering quantity, due originally to Shannon.*

Chapter 8 introduced admissibility volume through bandwidth. The Shannon formula suggests an extension.

Definition 33 (Informational admissibility volume). *The informational admissibility volume of a communication channel is the set of distinguishable signal configurations that remain recoverable in the presence of the channel's noise environment.*

Observation 106. *Bandwidth alone does not determine this volume; noise constrains it as well. A channel with enormous bandwidth but overwhelming noise may admit fewer reliable distinctions than a narrower but cleaner channel. This is one of the clearest places where admissibility becomes quantitative rather than merely conceptual.*

21.6 The Compactness Criterion

Definition 34 (Physical compactness criterion). *A signal representation is physically compact if it preserves the distinctions necessary for successful recovery while minimizing the resources required for transmission, storage, or processing.*

Observation 107. *Compression and communication are always constrained by noise. A representation that is too compact may become unrecoverable once fluctuations are introduced; a representation that is too redundant wastes resources. Engineering practice negotiates continually between compactness and recoverability: source coding removes redundancy, error-correcting codes reintroduce carefully chosen redundancy, and both are responses to the same constraint.*

21.7 A Circuit-Theoretic Example

Consider a digital communication link transmitting binary symbols, with logical 0 corresponding to voltage V_0 and logical 1 to voltage V_1 , both perturbed by noise.

Observation 108. *As long as the separation between V_0 and V_1 remains large relative to the noise amplitude, the receiver reliably distinguishes the two states. As noise increases, the overlap between the distributions grows until the states become difficult, and eventually impossible, to distinguish with acceptable reliability. The failure is not energetic; both voltages still exist. The failure is distinguishability-based: the receiver can no longer tell them apart. This is the practical meaning of low signal-to-noise ratio.*

21.8 Relation to Earlier Chapters

Observation 109. *Chapter 8's admissibility gates determine which frequencies are permitted to continue. Chapter 9's diode introduces irreversible collapse. Chapter 20's entropy production introduces fluctuation. Signal-to-noise ratio sits at the intersection of all three: admissibility determines what may pass, irreversibility determines what may be lost, and fluctuation determines what remains distinguishable after transmission. This is why SNR occupies such a central position in engineering practice despite its simple definition: it measures the survival of distinction under noise.*

21.9 Where the Compactness Picture Breaks Down

Observation 110. *Signal-to-noise ratio is fundamentally a second-moment quantity. Two systems with identical SNR may possess dramatically different error characteristics if their noise distributions differ substantially from one another. Many modern communication systems operate near limits where coding structure, algorithmic complexity, and statistical assumptions matter as much as raw SNR. This chapter should be read as a first-order interpretation, not a complete theory of communication.*

21.10 Verdict

Chapter 20 established that noise is an unavoidable consequence of fluctuation and dissipation. This chapter identifies what noise threatens: distinction. Signal-to-noise ratio is more than a power ratio; it is a measure of recoverable structure. Bandwidth determines how much frequency space is available; noise determines how much of that space remains distinguishable; Shannon capacity combines both into a single quantitative limit. Signal-to-noise ratio may therefore be understood as a compactness criterion imposed by physics itself: any representation, communication protocol, or measurement procedure must ultimately preserve enough distinction to survive the fluctuations of the world through which it passes.

22 Failure Modes and Limitations

A framework that only accumulates supporting examples is not yet a theory. It is a vocabulary. The earlier chapters of this book have repeatedly attempted to avoid this failure by distinguishing between exact correspondences, useful reinterpretations, and speculative extensions. Chapter 5 explicitly asked whether the amplitwist language did real conceptual work rather than merely renaming familiar quantities. Chapter 6 identified the boundary beyond which scalar amplitwist composition gives way to matrix composition. Chapter 9 limited its discussion of thermodynamic erasure. Chapter 20 stated explicitly where the entropy-production picture becomes incomplete.

This chapter collects those limitations into a single place and adds several new ones. The purpose is not defensive. The purpose is the same one served by the falsification chapter of the locomotion volume: a framework should be capable of being wrong.

22.1 What This Book Actually Claims

Observation 111. *The strongest claims of this book are narrow. Impedance is literally an amplitwist operator. Matched transmission coefficients compose exactly like amplitwist operators. Injection locking and Kuramoto synchronization share the same phase-coupling equations. Capacitors and inductors are literally history-dependent devices. Johnson–Nyquist noise is quantitatively tied to dissipation through the fluctuation-dissipation theorem. These claims are either definitional or independently established in existing engineering literature.*

Observation 112. *The weaker claims concern interpretation. The language of admissibility, continuation, amplitwist fixed points, and history-primacy is intended as a unifying description of phenomena already known to circuit theory. The burden of proof therefore falls not on whether the phenomena exist, but on whether the unified description explains anything that ordinary circuit language leaves obscure.*

This distinction matters because different claims fail in different ways.

22.2 Failure Mode I: Idle Renaming

Definition 35 (Idle-renaming failure). *A reinterpretation suffers idle-renaming failure if every statement expressible in the reinterpreted language is merely a restatement of existing theory, with no gain in explanatory power, compression, organization, prediction, or conceptual clarity.*

Observation 113. *This is the central danger facing the amplitwist vocabulary. Since impedance is already represented by complex multiplication, calling it an amplitwist operator is trivially true. The question is whether that truth accomplishes anything. The pilot chapter (Chapter 5) identified one candidate benefit: distinguishing operator addition from operator composition. Whether that benefit is substantial enough to justify the vocabulary remains an open question. If no further benefits emerge elsewhere in the book, the amplitwist interpretation should be regarded as pedagogically optional rather than theoretically important.*

22.3 Failure Mode II: Breakdown of Scalar Amplitwist Structure

Observation 114. *A scalar amplitwist operator is a single complex number acting by multiplication. Many real systems are more complicated.*

Proposition 25. *A mismatched cascade of two-port networks generally cannot be represented by multiplication of scalar transmission coefficients alone.*

Proof. Reflections introduce multiple internal propagation paths. The resulting behavior depends on the full network matrix rather than on a single transmission coefficient. \square

Observation 115. *This observation, established in Chapter 6, generalizes: distributed microwave structures, cavity resonators, multiport networks, and coupled electromagnetic fields may require matrix-valued or operator-valued descriptions whose behavior cannot be compressed into a single amplitude and phase pair. The amplitwist language survives only by generalizing beyond the scalar case or by accepting a restricted domain of applicability.*

22.4 Failure Mode III: Non-Sinusoidal Regimes

Much of the mathematics developed in this book depends on phasor analysis, and phasor analysis depends on sinusoidal steady state.

Definition 36 (Non-sinusoidal regime). *A circuit operates in a non-sinusoidal regime when no single-frequency phasor adequately describes the dominant system behavior. Examples include switching regulators, pulse-width modulation, digital logic transitions, relaxation oscillators, and strongly nonlinear waveforms.*

Observation 116. *In such systems a Fourier decomposition remains possible, but the amplitwist interpretation becomes distributed across many harmonics rather than concentrated in a single operator. Whether a useful higher-level amplitwist description exists is presently unknown. This is the reason the switching-converter problem was flagged early in this book as a likely stress test.*

22.5 Failure Mode IV: Adaptive Admissibility

Chapter 8 interpreted filters as admissibility gates. That interpretation works cleanly only when the admissibility region is fixed.

Observation 117. *Adaptive filters alter their behavior in response to incoming signals; the admissibility region therefore becomes history-dependent rather than fixed. The difficulty is not that admissibility disappears, but that it becomes a dynamical object. The simple geometric picture developed in Chapter 8 may require substantial extension before it can describe modern adaptive systems.*

22.6 Failure Mode V: Lock Loss and Synchronization Collapse

The Synchronization and Phase-Locked Loop chapters rely heavily on coherence and phase locking. Those concepts have limits.

Proposition 26. *The order parameter loses explanatory value when a population fragments into multiple coherent clusters rather than remaining synchronized or fully incoherent.*

Proof. The order parameter compresses a population into a single coherence magnitude and mean phase. Multiple coherent clusters can produce order-parameter values indistinguishable from those of qualitatively different synchronization structures, since the compression discards cluster-level detail by construction. \square

Observation 118. *This limitation is already known in the synchronization literature independent of this book, most concretely in the study of chimera states, in which a population of identically coupled oscillators spontaneously splits into a coherent and an incoherent subpopulation (Kuramoto and Battogtokh, 2002; Abrams and Strogatz, 2004). The amplitwist reading inherits this limitation because it inherits the order parameter itself. The same issue appears in PLLs operating near cycle slips, lock loss, or strongly nonlinear acquisition regimes, where small-signal phase descriptions cease to be sufficient.*

22.7 Failure Mode VI: Entropy Without Dissipation

Observation 119. *Chapter 20 linked noise and entropy production through the fluctuation-dissipation theorem, but the connection, while powerful, is incomplete: shot noise, flicker*

noise, chaotic dynamics, and informational entropy need not reduce cleanly to the resistor-based entropy production framework developed there. This limitation is particularly important because the book repeatedly invokes entropy as part of the RSVP triad; a successful theory of entropy for circuits must eventually account for more than thermal dissipation alone.

22.8 Failure Mode VII: History Without Continuation

Several chapters argue that capacitors, inductors, feedback loops, and PLLs are fundamentally historical systems. That observation can also be overextended.

Observation 120. *Not every history-dependent system is a continuation process in the sense developed elsewhere in this author's work. A device may possess memory without actively regenerating the conditions of its own persistence. This distinction matters because otherwise every stateful system would automatically count as a continuation system, rendering the concept too broad to be useful. History is necessary for continuation; it is not sufficient.*

22.9 A Compactness Criterion for the Framework Itself

The locomotion volume proposed a compactness criterion: an operator description should not require more complexity than the trajectory description it replaces. The same principle applies here.

Definition 37 (Conceptual compactness criterion). *The continuation geometry framework satisfies conceptual compactness if the explanatory machinery it introduces remains smaller than the collection of disconnected explanations it replaces.*

Observation 121. *If amplitwist, admissibility, continuation, synchronization, and history-primacy ultimately require more concepts than ordinary circuit theory already possesses, the framework has failed by its own standards. A unifying language that expands complexity rather than compressing it is not functioning as a unifying language. This criterion is intentionally severe: a successful reinterpretation should reduce conceptual burden rather than increase it.*

22.10 Summary

The framework developed in this book can fail in at least seven distinct ways: it can collapse into idle renaming; it can fail outside scalar amplitwist regimes; it can fail in strongly non-sinusoidal systems; it can fail when admissibility regions become adaptive; it can fail when synchronization fragments into chimera-like states; it can fail to account for important forms of entropy and noise; and it can fail by confusing

memory with continuation. Most importantly, it can fail the compactness criterion and become more complicated than the theory it seeks to unify.

These possibilities are not embarrassments. They are conditions under which the framework becomes informative. A theory that cannot fail cannot succeed either. For that reason the limitations collected here should be read not as exceptions to the book's argument but as part of its argument: the claim that continuation geometry is useful acquires meaning only because there are identifiable situations in which it may cease to be.

23 Two-Port Cascades and Assembly Index: A Comparative Note

This chapter is more modest than most of what has preceded it, and it is worth saying so directly rather than manufacturing a load-bearing result where none is needed. Chapter 6 already defined an Assembly Index for matched two-port cascades; this book's Chapter 4 recalled a different Assembly Index, from this author's locomotion volume, for hierarchical amplitwist chains. Both are called Assembly Index. This chapter's only task is to check whether that shared name hides a shared structure or merely a coincidence, in the same spirit as Chapter 7's disambiguation of three different senses of "fixed point."

23.1 Two Different Composition Operations, One Shared Name

Observation 122. *Chapter 6's Assembly Index counts sequential compositions: each additional cascaded two-port stage multiplies one more transmission coefficient into the running product, $A(\text{cascade}) = n - 1$ for n stages. The locomotion volume's Assembly Index counts hierarchical compositions: each additional level applies a coupling functional C that combines a population of lower-level operators into one coarser-level operator. These are not the same operation performed in two settings; multiplication of two scalars and coupling-and-coarsening of a population are structurally different acts, related only in that both are, in some sense, "one more join."*

Definition 38 (General Assembly Index, recalled). *This author's prior work on procedural ontology defines, for a primitive basis B and a fixed set of admissible composition rules, $A(P) = \min\{n : P \in \text{Span}_n(B)\}$, the minimum number of irreducible joins needed to construct P from B .*

Observation 123. *Both of the specific Assembly Indices used elsewhere in this book are instances of Definition 1, differing only in what counts as a primitive basis element and what counts as an admissible join. For Chapter 6, B is a set of individual matched two-port stages and the join is complex multiplication of transmission coefficients. For the locomotion volume's chains, B is a population of finest-scale amplitwist operators and the join is*

the coupling functional C . Definition 1 is genuinely general enough to cover both without strain, which is a modest but real clarification: the two Assembly Indices are siblings under one schema, not two unrelated uses of a convenient word, and not secretly identical either.

23.2 A Natural, Unresolved Extension

Observation 124. *Nothing in this book has considered a mixed case: a cascade of stages, each of which is itself a coupled population rather than a single element, for instance a sequence of coupled-oscillator-array stages connected one after another. Such a system would have joins of both kinds, sequential multiplication between stages and hierarchical coupling within each stage, and its Assembly Index, under Definition 1, would presumably depend on both the number of cascaded stages and the population size within each. This book does not work out this case; it is noted here as a natural next question the comparative framing of this section makes visible, not as a result this chapter has established.*

23.3 Verdict

This chapter's contribution is narrower than most others in this book, and its own verdict should say so plainly: it does not introduce new circuit behavior, and its one genuine result, that Chapter 6's and the locomotion volume's Assembly Indices are both instances of one general schema, is a taxonomic clarification rather than a load-bearing theorem. It is worth having stated once, given that the same word has now been used twice in this author's work with two different underlying operations, but it does not carry the weight of, for instance, Chapter 9's rectification theorem or Chapter 17's information-theoretic gate comparison.

24 Open Problems

This closing chapter collects the specific open questions flagged, honestly, throughout the rest of this book, rather than introducing a final new unifying claim. Chapter 22 already stated the conditions under which this book's central framework could fail; this chapter instead assumes the framework has, for now, survived those conditions, and asks what remains to be worked out.

24.1 The Full Matrix Case

Observation 125. *Chapter 6 showed that scalar amplitwist composition holds exactly only for matched cascades, and that mismatched cascades require the full ABCD matrix treatment. Whether a higher-dimensional generalization of the amplitwist decomposition, some structured decomposition of a general 2×2 matrix into an amplitude-and-twist-like normal form, could recover a comparably clean geometric reading of the mismatched case is left*

open. This book has treated the matched case as the exact special case worth understanding deeply, not as a stepping stone to a general theory it does not yet have.

24.2 Nonequilibrium Entropy and the Diode

Observation 126. *Chapter 9 declined to invoke Landauer's bound for continuous diode conduction, on the grounds that the bound concerns discrete, clocked erasure and a diode is neither. Whether a nonequilibrium statistical mechanics treatment could justify an analogous bound for continuous, unclocked dissipative processes is a real question in statistical physics independent of this book, and this book's own contribution was to identify precisely where the boundary of the simpler, equilibrium-flavored argument sits, not to resolve what lies beyond it.*

24.3 Non-Periodic and Transient Behavior

Observation 127. *This book's oscillator and synchronization chapters depend on populations of periodic phase oscillators. Ordinary transient circuit behavior, power-up sequences, single-shot pulse circuits, one-time switching events, is not periodic in this sense, and this book has not asked whether a degenerate, single-event amplitwist reading (in the spirit of the locomotion volume's own open question about non-periodic movement) has anything useful to say about transients. This is left for future work rather than attempted here.*

24.4 Reversible and Adiabatic Circuits

Observation 128. *Chapter 17 noted that reversible logic gates (Toffoli, Fredkin) carry no Landauer-bound dissipation in principle, since they are bijective. Adiabatic or reversible CMOS circuits, an active area of research independent of this book (see, e.g., Koller and Athas), attempt to approach this zero-dissipation limit physically, by charging and discharging capacitances gradually rather than abruptly. Whether such circuits are best understood, in this book's vocabulary, as physical realizations of the invertible, amplitwist-like side of the reversible/irreversible split rather than as an engineering trick unrelated to it, is a natural question this book raises but does not pursue.*

24.5 Correlated, Multi-Cell Failure and Repair

Observation 129. *Chapter 19 treated DRAM refresh and single-bit error correction under an assumption of independent, per-cell leakage and per-bit error. Real DRAM exhibits correlated, multi-cell disturbance effects, most notably row-hammer-style disturbance errors, in which repeatedly accessing one memory row can induce bit flips in physically adjacent rows (documented by Kim et al.). Whether this book's repair vocabulary, developed for independent per-bit or per-cell failure, extends cleanly to correlated, spatially structured failure of this kind, or requires a substantially different formal treatment, is left open.*

24.6 Admissibility as a Design Methodology

Observation 130. *This book has used admissibility throughout as an analytical tool: given a circuit, determine whether some configuration or operating condition is admissible relative to a stated criterion. A different and more ambitious use would be as a synthesis tool: given an admissibility criterion, construct a circuit guaranteed to satisfy it, working forward from the criterion rather than backward from an existing design. This book has not attempted this, and it is noted here as a direction rather than a result.*

24.7 Beyond Circuits

Observation 131. *This book and the locomotion volume it draws vocabulary from both apply a broadly similar continuation-geometric framework to two specific physical domains. Whether the same framework has anything to say about chemical reaction networks, fluid systems, or other physical continua entirely is a genuinely open question this book does not attempt to answer, and should not be assumed to have an affirmative answer merely because the framework has now been applied twice. The discipline this book has tried to maintain throughout, stating falsification conditions before claiming success, applies with at least as much force to any future extension beyond circuits and locomotion as it did within them.*

Appendix: A Classification of This Book's Claims

Every chapter of this book has closed with a verdict distinguishing what it actually established from what it merely redescribed. This appendix makes that recurring distinction explicit and systematic, sorting the book's major claims into three tiers:

Tier 1 — Standard circuit-theoretic facts. Results already fully established in ordinary circuit theory or physics, independent of anything in this book, stated here without modification.

Tier 2 — Structural equivalences proven in this book. Claims this book derives, with a proof or a decisive citation to an independent convergence, that would hold whether or not the reader accepts any of the book's broader vocabulary.

Tier 3 — Interpretive or organizational proposals. Choices of vocabulary, framing, or classification that are useful for organizing the material and connecting it to this author's broader work, but that are not themselves discovered mathematical structure, and should not be mistaken for it.

The purpose of this classification is not to demote Tier 3. Organizing vocabulary has genuine value, and several Tier 3 proposals (admissibility, history-before-state, Pop/Refuse) recur usefully across many chapters. The purpose is to prevent exactly the failure this book has tried to avoid throughout: allowing a reader to lose track of which claims would survive a skeptic rejecting the book's framework entirely, and which would not.

Foundations (Chapters 1–4)

Tier 1: Voltage as electric scalar potential Φ ; KCL as charge conservation; KVL as the zero-circulation property of a conservative field (quasi-static limit). **Tier 2:** Admissibility can be empty (Proposition 1, Chapter 4) — a genuine, proven boundary condition. **Tier 3:** The general physical/contextual admissibility framework itself; the memoryless/historical constraint taxonomy, offered as organizing vocabulary for later chapters rather than as new physics.

Phasors as Amplitwist Operators (Chapter 5)

Tier 1: $Z = |Z|e^{j\theta}$; Ohm’s law; series and parallel combination formulas. **Tier 2:** None. This chapter’s own verdict states plainly that it produces no new predictions. **Tier 3:** The identification of impedance as an “amplitwist operator”; the operator-addition-versus-composition distinction (a real conceptual clarification, but a clarification of vocabulary, not a new provable fact); the resonance-as-twist-vanishing reframing, explicitly flagged in the chapter itself as unproven pedagogical conjecture.

Two-Port Networks as Amplitwist Chains (Chapter 6)

Tier 1: ABCD matrix composition under cascading. **Tier 2:** Matched-cascade transmission coefficients compose exactly like amplitwist operators, amplitude multiplying and phase adding (Proposition 2) — a genuine derived result, not a relabeling, since it identifies a specific algebraic law (the chain rule for composed holomorphic derivatives) as the exact mechanism. **Tier 3:** The “amplitwist chain” vocabulary; the Assembly Index label applied to cascade depth.

Resonance as an Amplitwist Fixed Point (Chapter 7)

Tier 1: LC resonance condition $\omega_0 = 1/\sqrt{LC}$; reactance cancellation. **Tier 2:** None. Proposition 1 (resonance iff zero twist) is a direct restatement of the resonance condition in amplitwist vocabulary, not an independent derivation. **Tier 3:** The entire “amplitwist fixed point” framing, explicitly distinguished in this chapter from the unrelated dynamical fixed point of Chapter 13; the resonance/synchronization connection, explicitly marked in the chapter as qualitative and underived.

Filters as Admissibility Gates (Chapter 8)

Tier 1: Passband, stopband, and transfer-function definitions; bandwidth as a frequency-domain measure. **Tier 2:** None. This is, by the book’s own reckoning and the critique that prompted this appendix, the chapter where the case for discovered

structure is weakest. **Tier 3:** The entire admissibility-gate vocabulary applied to filtering. A reader may reasonably conclude, as the critique motivating this appendix suggests, that admissibility is a useful lens on filtering rather than a structure filtering theory independently requires.

The Diode as an Irreversible Spheredpop Operator (Chapter 9)

Tier 1: The Shockley diode equation; forward/reverse bias behavior. **Tier 2:** Rectification requires a non-invertible operator (Theorem 1) — an elementary but fully independent proof, requiring no acceptance of amplitwist, Spheredpop, or continuation vocabulary to be true or useful. **Tier 3:** The Pop/Refuse naming; the explicit, deliberate decision *not* to invoke Landauer’s bound here (itself a Tier-3-adjacent methodological choice, not a claim).

Histories Before States (unnumbered chapter, after Chapter 9)

Tier 1: Capacitor and inductor constitutive relations in integral form. **Tier 2:** Capacitor voltage is a non-injective functional of current history (Proposition 1) — a genuine proof, and the observation that standard circuit theory already calls v_C, i_L “state variables” for exactly this reason, which is an independently verifiable fact about existing terminology, not an interpretive overlay. **Tier 3:** The “history before state” framing as a general thesis; the extension of this idea, by named connection rather than proof, to diode reverse recovery and the PLL integrator.

Active Devices and Analog Computation (Chapters 10–12)

Tier 1: BJT/FET region equations; op-amp golden rules; the log-amplifier and log-antilog multiplier circuits. **Tier 2:** The active-region/cutoff-saturation classification of BJT invertibility (a genuine, checkable regional distinction); the ideal-op-amp admissibility gap as an exact function of A_{ol} (Chapter 11); logarithmic compression of dynamic range (Proposition 1, Chapter 12) — a genuine, quantitative, provable claim; the correction of the log-multiplier derivation to the standard log-antilog technique, which is simply correct engineering, independent of this book. **Tier 3:** The “admissibility engine” framing of op-amp analysis; the Assembly-Theoretic compression analogy, explicitly marked in the chapter as structural/quantitative rather than operational.

Feedback and the Barkhausen Criterion (Chapters 13–14)

Tier 1: The closed-loop gain formula $X_o = \frac{A}{1+A\beta} X_i$; the Barkhausen magnitude and phase conditions. **Tier 2:** The Barkhausen criterion is exactly the linear onset condition for a Hopf bifurcation (Proposition 1) — a genuine correspondence between

two independently developed theories; the necessary-but-not-sufficient distinction for self-starting oscillation, a real correction to a common oversimplified statement of Barkhausen’s criterion. **Tier 3:** Reading the closed-loop equation as “recursive continuation.” This is the one Tier 3 item in the book closest to requiring no interpretive leap at all, since the fixed-point structure is literally present in the algebra; it is listed here rather than in Tier 2 because the algebra alone does not require the vocabulary of continuation to be understood or used correctly.

Oscillator Circuits (Chapter 15)

Tier 1: RC phase-shift, Hartley, and crystal oscillator frequency formulas (the RC result cited rather than re-derived). **Tier 2:** None beyond the direct application of Chapters 13–14’s already-established Tier 2 results to specific topologies. **Tier 3:** None especially new; this chapter is framed throughout as a worked instantiation of earlier material rather than a source of new claims.

Synchronization

Tier 1: None; Adler’s equation and the Kuramoto model are themselves independently established results in electrical engineering and statistical physics, predating this book by decades. **Tier 2:** The identification of Adler’s equation with the two-oscillator Kuramoto equation is Tier 2 in the strongest sense available anywhere in this book: it is a documented historical convergence between two literatures that never referenced each other, not a construction of this book’s own. **Tier 3:** The admissibility framing of lock and synchronization.

The Phase-Locked Loop

Tier 1: The PD/LF/VCO block-diagram equations. **Tier 2:** A Type-I PLL is exactly Adler’s equation with an engineered coupling constant (Proposition 1); a Type-II PLL achieves zero steady-state phase error under a frequency step (Proposition 2) — both genuine derived results with stated domains of validity. **Tier 3:** The “engineered continuation” framing of the chapter’s title.

Digital Logic as Spherepop Collapse (Chapter 17)

Tier 1: Standard logic gate truth tables. **Tier 2:** Every two-input gate is non-invertible (Proposition 1, elementary but real); AND and XOR destroy quantitatively different amounts of information (Proposition 2); the resulting Landauer dissipation floor (Proposition 3) — a bound properly earned in this discrete, clocked setting where Chapter 9 correctly declined it for a continuous device. **Tier 3:** The Pop/Refuse vocabulary applied to gates; the metastability-as-Chapter-13-fixed-point

connection, which is a genuine physical fact about flip-flop circuits but is presented as an illuminating instance of earlier vocabulary rather than a new proof.

Clock Distribution, Skew, and Domain Crossing (Chapter 18)

Tier 1: Setup/hold timing inequalities; two-flop synchronizer practice. **Tier 2:** Skew-induced timing admissibility can be empty (Proposition 1) — a second, independent instance of Chapter 4’s admissibility-emptiness result, in a genuinely different domain (digital timing closure rather than oscillator design). **Tier 3:** The general admissibility vocabulary applied to clock timing.

Memory and Repair (Chapter 19)

Tier 1: DRAM leakage and refresh; Hamming(7,4) encoding and syndrome decoding. **Tier 2:** None strictly new mathematically; the refresh-interval inequality and the Hamming syndrome mechanics are standard results, correctly stated. **Tier 3:** The mapping of Hamming decoding onto this author’s Validate/Repair formalism (a genuine structural match, but organizational rather than a new theorem); the “silent miscorrection” framing, which names a real and important phenomenon but does so descriptively rather than by proof; DRAM refresh as “recursive continuation.”

Noise as Entropic Production (Chapter 20)

Tier 1: Johnson–Nyquist noise; shot noise; the fluctuation-dissipation theorem itself, established in statistical mechanics independent of this book. **Tier 2:** None beyond correctly reporting Tier 1 results; this chapter’s contribution is connective rather than derivative. **Tier 3:** Reading dissipation as entropy production within this book’s own (Φ, v, S) vocabulary; the parallel drawn to the “Second Law for Code” of this author’s other work, explicitly marked as a separate, independent instance rather than a shared derivation.

Signal-to-Noise Ratio (Chapter 21)

Tier 1: SNR definitions; Shannon’s channel capacity theorem, cited rather than derived. **Tier 2:** None; this chapter is explicitly interpretive throughout. **Tier 3:** The recoverability and distinguishability framing of SNR; “informational admissibility volume”; the physical compactness criterion.

Failure Modes and Limitations (Chapter 22)

This chapter is, by design, entirely a synthesis of limitations already stated in Tier 2 and Tier 3 items above, together with the book's own compactness criterion. It does not introduce new Tier 1 or Tier 2 claims; its function is to state, in one place, the conditions under which every interpretive proposal elsewhere in the book could be judged to have failed.

Two-Port Cascades and Assembly Index; Open Problems (Chapters 23–24)

Tier 1: None. **Tier 2:** The observation that this book's two distinct Assembly Indices (sequential-multiplicative and hierarchical-additive) are both instances of one general schema is a modest but genuine clarification, not a new physical result. **Tier 3:** Everything in the Open Problems chapter, by its own declared nature, including its explicit refusal to assume the book's framework generalizes beyond circuits and locomotion merely because it has now been applied twice.

What This Classification Does Not Resolve

Sorting claims into three tiers does not, by itself, settle the harder question this appendix's own motivating critique raised: whether the Tier 3 proposals recurring across chapters (admissibility, history-before-state, Pop/Refuse) are compatible reinterpretations of independently true facts, or symptoms of one deeper organizing principle circuit theory actually possesses. This book has not proven the latter, and this appendix does not attempt to. What the classification does is make it possible to see, chapter by chapter, exactly how much of the book's force depends on which claim, so that a reader inclined to accept the Tier 2 results while remaining skeptical of the Tier 3 unification can do so consistently, without either underrating what has been proven or overrating what has only been proposed.

References

- [1] C. A. Desoer and E. S. Kuh (1969). *Basic Circuit Theory*. McGraw-Hill.
- [2] D. M. Pozar (2011). *Microwave Engineering*, 4th ed. Wiley.
- [3] P. Horowitz and W. Hill (2015). *The Art of Electronics*, 3rd ed. Cambridge University Press.
- [4] A. S. Sedra and K. C. Smith (2014). *Microelectronic Circuits*, 7th ed. Oxford University Press.

- [5] R. L. Boylestad and L. Nashelsky (2013). *Electronic Devices and Circuit Theory*, 11th ed. Pearson.
- [6] R. Adler (1946). "A Study of Locking Phenomena in Oscillators." *Proceedings of the IRE* 34(6): 351–357.
- [7] R. A. York and R. C. Compton (1991). "Quasi-optical power combining using mutually synchronized oscillator arrays." *IEEE Transactions on Microwave Theory and Techniques* 39(6): 1000–1009.
- [8] K. Kurokawa (1973). "Injection locking of microwave solid-state oscillators." *Proceedings of the IEEE* 61(10): 1386–1410.
- [9] R. E. Best (2007). *Phase-Locked Loops: Design, Simulation, and Applications*, 6th ed. McGraw-Hill.
- [10] R. Landauer (1961). "Irreversibility and Heat Generation in the Computing Process." *IBM Journal of Research and Development* 5(3): 183–191.
- [11] C. H. Bennett (1973). "Logical Reversibility of Computation." *IBM Journal of Research and Development* 17(6): 525–532.
- [12] S. Harris and D. Harris (2021). *Digital Design and Computer Architecture*, 2nd ed. Morgan Kaufmann.
- [13] T. J. Chaney and C. E. Molnar (1973). "Anomalous Behavior of Synchronizer and Arbiter Circuits." *IEEE Transactions on Computers* C-22(4): 421–422.
- [14] R. W. Hamming (1950). "Error Detecting and Error Correcting Codes." *Bell System Technical Journal* 29(2): 147–160.
- [15] J. B. Johnson (1928). "Thermal Agitation of Electricity in Conductors." *Physical Review* 32: 97–109.
- [16] H. Nyquist (1928). "Thermal Agitation of Electric Charge in Conductors." *Physical Review* 32: 110–113.
- [17] W. Schottky (1918). "Über spontane Stromschwankungen in verschiedenen Elektrizitätsleitern." *Annalen der Physik* 362(23): 541–567.
- [18] C. E. Shannon (1948). "A Mathematical Theory of Communication." *Bell System Technical Journal* 27(3): 379–423.
- [19] Y. Kuramoto and D. Battogtokh (2002). "Coexistence of Coherence and Incoherence in Nonlocally Coupled Phase Oscillators." *Nonlinear Phenomena in Complex Systems* 5(4): 380–385.

- [20] D. M. Abrams and S. H. Strogatz (2004). “Chimera States for Coupled Oscillators.” *Physical Review Letters* 93(17): 174102.
- [21] J. G. Koller and W. C. Athas (1992). “Adiabatic Switching, Low Energy Computing, and the Physics of Storing and Erasing Information.” *Proceedings of the Workshop on Physics and Computation*: 267–270.
- [22] Y. Kim, R. Daly, J. Kim, C. Fallin, J. H. Lee, D. Lee, C. Wilkerson, K. Lai, and O. Mutlu (2014). “Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors.” *ACM SIGARCH Computer Architecture News* 42(3): 361–372.