Synchrony and Release

Hierarchical Oscillator Theory with Pneumatic Paper Robots

Flyxion
Independent Researcher

November 2026

Abstract

This monograph develops a theory of coordination for robots built from oscillators arranged in a hierarchy, and applies it to a proposed class of machines: lightweight paper-shelled robots whose limbs are pulled by internal cords drawn from a central spine and whose body can also be moved by air pressure distributed through gated chambers. The theory treats synchronization, selective release and recapture of coupled oscillators as separate operations on a coupling structure. Within a stated phase model it proves that a group of oscillators keeps its internal coordination within an explicit bound while other groups are released and recaptured, that the prescribed collective pattern is recovered with explicit error for two groups, for a three-level hierarchy and for three groups under stated conditions, and that these results survive adaptation of the internal couplings that stays within a certified box. For the pneumatic chambers, the unaveraged phase equation of a single-pressure chamber is derived exactly within the stated model (ideal threshold switching, linear gates with constant delays, and a bound on the gate inflow, Assumption W of Chapter 34); the averaged interaction equation used to connect the chambers to the phase model is an approximation to it, whose error is bounded only by a finite-horizon estimate with constants too large to be informative at the couplings used.

The pneumatic embodiment supplies a physical architecture for the abstract operations. A first-order derivation from chamber dynamics gives the interaction functions of three gate mechanisms, a worked three-chamber gait network is analyzed and simulated in two pneumatic cycles, and an interface carries chamber phases to cord lengths with an explicit error budget. Simulation evidence is limited to the model, one cycle needed recalibration to reproduce the behavior of the other, and the passage from cord lengths to joint motion and to balance is not addressed. The monograph also specifies an executable bench protocol for measuring the interaction function of two chambers, which has not been run. Each result is labeled by its scope: an exact identity, a theorem about the phase model or the exact chamber phase equation, a first-order approximation to the chamber dynamics, a numerical finding, or an untested physical proposal.

Scope of Results

Every result in this monograph belongs to one of five categories, and the categories are not interchangeable. An *identity* is an exact algebraic or definitional fact. A *theorem* is proved for the phase model of Part II, which describes oscillator phases and not the chamber dynamics that motivate it, and it holds under its stated hypotheses. An *approximation* is a first-order result derived from the chamber dynamics by averaging; the unaveraged phase equation of a single-pressure chamber is exact within its stated model and under Assumption W (Theorem 34.1), the averaged interaction equation obtained from it is not, and its error is bounded only by constants too large to be informative at the couplings used (Chapter 34). A *numerical finding* is a computed or simulated result at the parameters stated, with the step size and the number of runs as reported. A *proposal* is an untested physical design or hypothesis. The ledger of Appendix B assigns each obligation to a category and records its status, and each part begins by naming the categories that occur in it.

Two limits apply throughout. A theorem about the phase model does not by itself establish anything about a chamber or a body until the reduction from the physical system to the phase model has been bounded, and that bound is an open obligation. A numerical finding at one cycle does not establish behavior at another, as the second-cycle chapter shows.

Reading Guide

The monograph has two strands that can be read separately. The theory strand is Parts I to IV: oscillator definitions, coupling and hierarchy, the operations of maintenance, release, recapture and transition with their theorems, and the interpretation of summaries between levels. It is self-contained and does not need Part V. A reader interested only in the mathematics can read Chapter 2 only for the cycle used later, and then Parts II to III, with Part IV read as a short interpretive complement.

The embodiment strand is Chapter 2 of Part I together with Part V. A reader interested in the pneumatic construction can take the coordination results as statements to be quoted: Theorem 6.1 and Theorem 13.1 for internal cohesion, Theorem 7.1 for the effective oscillator, Theorem 9.1 for recapture, Theorem 13.3 for release and recapture of two groups, Theorem 14.1 for adaptation, and Theorems 25.1 to 25.3 for three groups. Their hypotheses are stated where they appear and are summarized in the scope note at the start of Part III. The embodiment strand then runs through the spine, the first embodiment, the interaction functions, the worked network and the second cycle, then the error of the averaged chamber model, and ends with the validation program and the bench protocol.

A third path concerns verification only: the validation program, the bench protocol, the error of the averaged chamber model, and the ledger. The ledger in Appendix B is the quickest way to see what is proved, what is approximated, what is simulated and what is untested. The Notation table at the end of the front matter lists the symbols and their uses, including the few that serve more than one purpose in different parts.

Preface

This monograph develops a theory of coordination built from coupled oscillators arranged in a hierarchy, and then applies it to a particular class of machines: lightweight robots whose bodies are hollow paper shells, whose limbs are pulled by internal cords drawn from a central spine, and whose body can also be moved by air pressure distributed through a network of cells and gates, independently of the cords, with liquid pressure as a proposed extension. These machines are called endomarionettes here. They are the application, and the theory of oscillators is the subject.

The central claim concerns what control consists of. Maintaining a coordinated movement means establishing phase relationships among oscillating subsystems. Acting independently means selectively relaxing some of those relationships while leaving the rest intact. Changing from one pattern of movement to another means establishing new relationships. Synchronization, release and transition are therefore treated as distinct operations on a coupling structure, and Part III is organized around that distinction.

A word on status is needed. Nothing in the application part has been built. The endomarionette is a design, described in earlier writing by the author, and the claims made about it here are claims about what the theory implies for such a body. Where a statement is a hypothesis, it is labeled as one and accompanied by the observations that would discriminate it from its alternatives. Where a result is a standard theorem, a reference is given.

The monograph works toward a single principal theorem, stated precisely in Part III: a hierarchically coupled oscillator system can selectively release a functional subsystem, preserve that subsystem's internal coordination, and subsequently recover a prescribed collective coordination pattern, with bounded error under specified perturbations. The endomarionette is then a physical realization whose dynamics must satisfy the assumptions of that theorem. The logical dependency runs from stable local oscillation, to stable group coordination, to justified coarse-graining, to selective release, to recapture and transition, to robust hierarchical control, and finally to physical realization. The geometric interpretation of Part IV branches from the coarse-graining results and interprets quantities justified along that chain; it does not carry the control argument. A ledger in Appendix B records, for each proof obligation, whether it is proved in this text, stated with a proof plan, or open.

Relation to Existing Work

Coupled central pattern generators for robot locomotion have substantial precedents, and this monograph does not claim them. Cheslet and colleagues implement a cascaded network of central pattern generators on a field programmable gate array for a wheeled snake-like

robot [1]. Each local generator consists of two mutually inhibitory neurons, and excitatory connections between generators introduce phase differences along the body. This is a hardware precedent for local rhythm generators joined into a larger movement pattern, with a forward cascade where this monograph uses internally cohesive groups. The same paper reports a preliminary closed-loop experiment between the artificial network and biological cells, performed once, in which the biological input modified the network but the robot's movement was disorganized, and it does not claim obstacle avoidance. Communication between components is not successful behavioral control, and the validation program of Part V keeps those apart.

Wang and colleagues give a controller for an ostraciiform fish robot with three functional layers: input saturation functions that transform the command inputs, coupled oscillators, and an output transition function into servomotor commands [22]. Two inputs regulate speed, gait and maneuvering, and gait changes work through command-dependent frequency and amplitude settings and through which fins contribute to propulsion. Their layers are processing stages and not nested levels of oscillator coarse-graining. Their gait changes also show that a change of movement need not proceed by weakening couplings, so release and recapture is one control operation alongside such established ones.

Duysens and Van de Crommert review the evidence for spinal rhythm generation in cats and humans, with sensory feedback, descending commands and phase-dependent reflex modulation, and they distinguish the rhythmic output of an isolated preparation from the full activation pattern of an animal walking in its environment [3]. Their evidence for a human generator is cited here as the assessment of a 1998 review and not as current consensus.

The work to be established in this monograph is narrower than these precedents: explicit preservation of internal coordination during release, bounded recapture, hierarchical error accounting, and the conditional realization of those operations through pneumatic connections. The sources support that positioning, and none of them validates the paper body or its mechanics.

Contents

I Single Oscillators **1**

1 Limit Cycles and Phase **3**
1.1 The Oscillator as a Reduced Object 3
1.2 Why Phase Is Enough for Coordination 3

2 Relaxation Oscillators and Their Pneumatic Realization **5**
2.1 The Pressure Cycle 5
2.2 From Chamber to Phase 6

3 Pneumatic Elements as Couplers **7**

II Coupling and Hierarchy **9**

4 Coupled Phase Oscillators **11**
4.1 The Basic Model 11
4.2 Prescribed Patterns and the Shift Lemma 11
4.3 The Pattern-Adjusted Order Parameter 12
4.4 Two Oscillators 13
4.5 Coordination Within a Group 13

5 Block-Structured Coupling **15**
5.1 Hierarchy as Structure of K 15
5.2 Why the Separation of Scales Matters 15

III Synchronization and Release **17**

6 Maintained Coordination **19**
6.1 Cohesion of a Group 19

7 The Effective Oscillator **23**
7.1 One-Level Reduction 23

8 Release **25**
8.1 Routes to Release 25

9 Independent Action and Recapture **27**
9.1 The Capture Theorem 27

10 Transition Between Coordination Patterns **29**
10.1 Four Candidate Mechanisms 29
10.2 Conditions 30
10.3 Discriminating Observations 30

11 Stopping, Amplitude and Engagement **31**

12 The Principal Theorem **33**

13 Nonuniform Coupling **35**
13.1 Setting 35
13.2 Cohesion Through the Internal Graph 35
13.3 Reduction and the Principal Theorem for Nonuniform Coupling 36
13.4 Amplitude and Offset Commands 37

14 Adaptation of Internal Couplings **39**
14.1 Admissible Region 39
14.2 A Law That Preserves the Region 39
14.3 The Adaptation Theorem 40
14.4 Moving Targets and Drifting Frequencies 41

15 Three-Level Hierarchies **43**
15.1 Setting 43
15.2 Propagation of the Remainders 43
15.3 Does the Error Estimate Stay Useful? 44

16 Boundary Simulations **45**
16.1 Nonuniform Two-Group Case 45
16.2 Initial Phases Near the Excluded Band 45
16.3 Asymmetric Forcing and Weak Internal Connections 45
16.4 Adaptation 46
16.5 Three Levels 47

IV Summaries Between Levels **49**

17 Summaries Passed Between Levels **51**
17.1 What the Results of Part III Supply 51
17.2 What the Summary Cannot See 52
17.3 Composition and Aggregation 52
17.4 The Open Interpretation 52

18 Chains of Joints, Limbs and Gaits **55**
18.1 Levels and Assembly Index 55
18.2 The Bipedal Case 55
18.3 Contact Events 55

V Endomarionettes **57**

19 The Endomarionette Body **59**
19.1 Three Mechanical Roles 59
19.2 The Blackbox 59
19.3 Embodiments of the Blackbox 60
19.4 Spinal Arrangements 60
19.5 Shared and Separate Take-Up 60

20 Cord and Pressure Actuation **61**
20.1 Routing Kinematics 61
20.2 Two Independent Channels 61

21 The Pneumatic Cell Network **63**
21.1 Cells, Gates and Reservoirs 63
21.2 A Linear Model 63
21.3 Release Events 64
21.4 Link to the Oscillator Theory 64

22 The Segmental Spine **65**
22.1 One Module 65
22.2 Connections Between Modules 66
22.3 Forming the Spine 66
22.4 A Directed Cascade Along the Spine 67

23 Channels, Constraints and Coupling Structure **69**
23.1 Shared Drums and Independent Channels 69
23.2 Which Channels Can Be Released 69
23.3 How Pressure Interacts with the Tendon Mechanism 70
23.4 The Resulting Coupling Structure 70

24 The First Embodiment **71**
24.1 The Decision 71
24.2 What Each Gate Changes 71
24.3 Residual Coupling 72

25 Worked Example: Stance, Free Leg and Arms **75**
25.1 The Three Groups 75
25.2 Preservation of Internal Coordination 75
25.3 Release and Recapture on a Pairwise Schedule 76
25.4 Simultaneous Recapture 76

25.5 A Simulation . . . 77
25.6 What the Theorem Does Not Establish . . . 78

26 Fabrication Routes **79**
26.1 Shell Forms . . . 79
26.2 Four Forming Concepts . . . 79
26.3 What Fabrication Must Deliver . . . 80

27 Loads on the Shell **81**
27.1 Guides . . . 81
27.2 Attachments . . . 81
27.3 Shell Stability . . . 81

28 Related Mechanisms and Precedents **83**
28.1 Paper and Origami Structures . . . 83
28.2 Tendon Transmission . . . 84
28.3 Pneumatic Muscles . . . 84

29 From Chamber Coupling to Phase Coupling **85**
29.1 Phase Response of the Pressure Cycle . . . 85
29.2 Weak Coupling Through a Gate . . . 86
29.3 Consequences for Design . . . 86

30 Interaction Functions and Gate Offsets **87**
30.1 The Interaction Function . . . 87
30.2 An In-Phase Connection . . . 88
30.3 A Half-Period Delayed Connection . . . 89
30.4 A Check-Valve Connection . . . 89
30.5 Sensitivity to Load . . . 90
30.6 Edges for a Gait Network . . . 90

31 From Chamber Phases to Cord Lengths **91**
31.1 The Mapping . . . 91
31.2 How the Take-Up Follows . . . 91

32 A Worked Three-Chamber Gait Network **93**
32.1 The Network . . . 93
32.2 Equations and the Locked Solution . . . 94
32.3 Margin After Loading . . . 94
32.4 Supply Sharing . . . 95
32.5 Release and Recapture . . . 96
32.6 Commanded Cord Lengths . . . 97

33 A Second Cycle **99**
33.1 The Two Cycles . . . 99
33.2 Interaction Functions . . . 99
33.3 The Network Under Four Settings . . . 100

33.4 The Retuned Network . . . 101
33.5 Release, Recapture and Drum Tracking . . . 102
33.6 What Is Retained and What Is Retuned . . . 103

34 Error of the Averaged Chamber Model 105
34.1 The Model and the Exact Phase Equation . . . 105
34.2 An Averaging Bound for the Exact Equation . . . 106
34.3 Loaded Chambers . . . 109
34.4 Diagnosis of the Cycle B Discrepancy . . . 111
34.5 What Is Established and What the Bound Must Cover . . . 115

35 Didactic Tasks as Coordination Sequences 117
35.1 Applying Paper Sheets . . . 117
35.2 Standing on One Leg While Dressing . . . 117
35.3 What the Examples Test . . . 118

36 Validation Program 119
36.1 Estimating Phase and Coherence . . . 119
36.2 Tests . . . 119

37 Bench Protocol for Two Chambers 123
37.1 Status and Purpose . . . 123
37.2 Quantities . . . 123
37.3 Normalization of the Slope Prediction . . . 123
37.4 Apparatus . . . 126
37.5 Stage 0: Calibration . . . 126
37.6 Stage 1: Each Chamber Alone . . . 127
37.7 Stage 2: The Uncoupled Pair . . . 127
37.8 Stage 3: Coupled Runs . . . 127
37.9 Estimating the Interaction Function . . . 127
37.10Fitting and Prediction Kept Apart . . . 129
37.11Linearity and Averaging Checks . . . 129
37.12The Comparator Simulation and Its Convergence . . . 129
37.13Reporting and Decision Rules . . . 133
37.14Later Stages and Limits . . . 133

38 Unresolved Design Choices 135

39 Conclusion 137

A Proposed Biological Extensions 139

B Ledger of Proof Obligations 141

Notation

Symbol	Meaning	Introduced
$\theta_i,\ \vartheta_i$	Phase of oscillator i; pattern-adjusted lifted phase $\theta_i - \beta_i$	Part I, II
$\omega_i,\ I_i$	Natural frequency; drive	Part II
$\beta_i,\ \alpha_{ij}$	Target offset of a pattern; preferred pairwise offset, $\alpha_{ij} = \beta_j - \beta_i$ for consistent patterns	Part II
$K_{ij},\ \kappa$	Coupling weights; uniform coupling strength	Part II
$Z_a = r_a e^{\mathrm{i}\Psi_a}$	Pattern-adjusted order parameter of group a, with argument Ψ_a	Part II
$\bar{\vartheta}_a$	Mean of the lifted phases of group a, distinct from Ψ_a	Part III
D_a	Spread of the lifted phases of group a	Part III
$W,\ E,\ \gamma$	Frequency spread, external-coupling bound, cohesion bound $\gamma = \arcsin((W+2E)/\kappa n)$	Part III
F	Bound on a sensory contribution (Corollary 6.1); also the label of the free-leg chamber in Part V	Part III, V
λ_a	Algebraic connectivity of the internal graph of group a	Part III
$\Phi_a,\ \Omega_a$	Disturbance norm of group a; Euclidean frequency dispersion	Part III
u	Relative phase between two groups, $\bar{\vartheta}_h - \bar{\vartheta}_g - a_{gh}$	Part III
$\mu(t)$	Control scaling the external coupling between groups	Part III
$N,\ \bar{N}$	Effective group-level coupling and its maximum	Part III
$s^*,\ \gamma_*$	Capture parameter and its arcsine	Part III
$y_i,\ a_i,\ e_i$	Output, amplitude and engagement of channel i	Part III
ξ_i	Sensory perturbation	Part III
$p,\ p_s,\ p_h,\ p_l,\ p_o$	Chamber pressure; supply, upper threshold, lower threshold, open-valve target	Part I
$\tau_c,\ \tau_o$	Charging and discharging time constants	Part I
T	Period of the pneumatic cycle (not cord tension)	Part I
$\mathcal{T}$	Cord tension	Part V
$Z(\theta)$	Phase response curve of the chamber to a pressure perturbation (distinct from Z_a)	Part V
$g,\ C,\ \epsilon = g/C$	Gate conductance, chamber capacitance, coupling rate	Part V
$G(\varphi),\ \bar{\Gamma}$	Averaged interaction functions of two chambers	Part V
$\epsilon' = \epsilon T_0,\ \Theta_i,\ \psi_i$	Coupling per cycle; lifted chamber phase; phase relative to the base rate	Part V
$\delta,\ \delta_F$	Relative natural rate difference of two chambers; extra detuning of the free-leg chamber	Part V
τ (delay)	Transport delay of a gate; distinct from the time constants τ_c, τ_o and the drum time constant	Part V
$S,\ F,\ A$	Stance, free-leg and arms chambers or groups	Part V
$\varsigma_a,\ L_s$	Bound on residual frequency perturbations; supply sensitivity	Part V

Part I

Single Oscillators

Chapter 1

Limit Cycles and Phase

Scope of this part. The definitions of phase and the phase response are standard and cited. The pneumatic period and cycle results of Chapter 2 are theorems about an idealized piecewise-linear chamber-and-valve model. The weak-forcing phase equation is a first-order approximation whose error is not bounded here, and the physical precedent for the valve is a literature reference, not a result.

1.1 The Oscillator as a Reduced Object

An oscillator, for the purposes of this monograph, is a dynamical system with a stable periodic orbit. Let $\dot{x} = f(x)$ on $\mathbb{R}^n$ have a stable limit cycle γ of period T and natural frequency $\omega = 2\pi/T$. Every point in the basin of γ can be assigned a phase $\theta \in \mathbb{T} = \mathbb{R}/2\pi\mathbb{Z}$ by the requirement that points on the same isochron converge to the same point of the cycle at the same time, so that $\dot{\theta} = \omega$ everywhere in the basin. This is the classical construction due to Winfree [24] and is developed in [14, 20]. A systematic modern account of the passage from a physical oscillator to a phase equation, covering isochrons, the phase sensitivity function, weak forcing and averaged interaction functions, is given by Nakao [13], and Chapter 29 uses its weak-coupling averaging.

The value of the construction is that a system of dimension n, perhaps large, is replaced by a single angle. For a weak perturbation $p(t)$ applied to the system, $\dot{x} = f(x) + p(t)$, the phase obeys

$$\dot{\theta} = \omega + Z(\theta) \cdot p(t), \tag{1.1}$$

where $Z(\theta) = \nabla_x \theta$ evaluated on the cycle is the infinitesimal phase response curve. Equation (1.1) holds to first order in the size of p. It will fail for strong forcing, and Part III returns to that limit when release and transition are discussed.

1.2 Why Phase Is Enough for Coordination

Coordination, as the term is used here, concerns timing relationships: which subsystem is at which point of its cycle when another is at a given point. Amplitude matters for what a movement looks like, but timing decides whether two movements are coordinated. Phase reduction discards amplitude information to isolate timing, and Chapter 11 restores amplitude as a separate variable, because the theory requires that the two can come apart.

Remark 1.1. The reduction to phase is an approximation with a stated domain. It assumes the cycle is strongly attracting relative to the perturbations applied. A paper body driven by pressure may violate this assumption near the limits of its range of motion, and the application chapters must check it rather than assume it.

Chapter 2

Relaxation Oscillators and Their Pneumatic Realization

2.1 The Pressure Cycle

Relaxation oscillators alternate slow accumulation with rapid discharge. The Van der Pol system in its strongly nonlinear regime is the standard example [21]. A pneumatic version requires no electronics, and this is the reason it is of interest for the machines treated later.

Consider a chamber of pressure $p(t)$ fed from a supply at pressure p_s through a restriction. A valve with hysteresis opens when p reaches an upper threshold p_h and closes when p falls to a lower threshold p_l, with $p_l < p_h < p_s$. While the valve is closed the chamber relaxes toward p_s with time constant τ_c. While the valve is open the chamber relaxes toward a lower target $p_o < p_l$ with time constant τ_o. This is an idealization: it treats both phases as first order and the switching as instantaneous.

Theorem 2.1. *Let $p_o < p_l < p_h < p_s$ and $\tau_c, \tau_o > 0$. Let the valve state $s \in \{0, 1\}$ (closed, open) switch from 0 to 1 when p reaches p_h and from 1 to 0 when p reaches p_l, with $\dot{p} = (p_s - p)/\tau_c$ when $s = 0$ and $\dot{p} = (p_o - p)/\tau_o$ when $s = 1$. Then from any initial state with $p \in [p_o, p_s]$ that is consistent with the switching rule, that is, with $p < p_h$ if $s = 0$ and $p > p_l$ if $s = 1$, the trajectory reaches a switching event in finite time and from then on is exactly periodic with period*

$$T = \tau_c \ln \frac{p_s - p_l}{p_s - p_h} + \tau_o \ln \frac{p_h - p_o}{p_l - p_o}. \tag{2.1}$$

On the cycle the two dwell times equal the two terms of (2.1), which are strictly positive, so switching events cannot accumulate in finite time. A state outside the consistent set, such as $s = 0$ with $p \geq p_h$, is not covered: with the closed valve it would stay at a rest point or decay without ever crossing p_h from below, and the model excludes it by the convention that the valve opens whenever $p \geq p_h$.

Proof. With the valve closed, p increases monotonically toward $p_s > p_h$ and so reaches p_h in finite time. With the valve open, p decreases monotonically toward $p_o < p_l$ and reaches p_l in finite time. Hence a switching event occurs after finite time, and the state just after it is $(p_h, 1)$ or $(p_l, 0)$, a point of the cycle. On the cycle, charging from p_l follows $p(t) = p_s - (p_s - p_l)e^{-t/\tau_c}$, which equals p_h at $t = \tau_c \ln \frac{p_s - p_l}{p_s - p_h}$. Discharging from p_h follows $p(t) = p_o + (p_h - p_o)e^{-t/\tau_o}$,

which equals p_l at $t = \tau_o \ln \frac{p_h - p_o}{p_l - p_o}$. Both logarithms are positive because $p_l < p_h$ and the targets lie outside the interval $[p_l, p_h]$. The trajectory after the first event is therefore the same closed curve regardless of the initial condition, which gives attraction in finite time, and the dwell times are the two terms. □

The hysteresis gap $p_h - p_l > 0$ is what excludes chatter: with $p_h = p_l$ both logarithms would vanish. This theorem settles the existence, attraction and well-posedness of the pneumatic cycle for the idealized model. It says nothing about leaks, finite valve switching time or delays, which belong to the robustness obligations listed in Appendix B.

A physical precedent. Rothemund and colleagues describe a soft elastomeric valve containing a bistable membrane that acts as a mechanical switch for air flow [15]. The snap-up and snap-down pressures of the membrane differ and can be designed through its geometry and material, a hysteresis of the kind the thresholds p_h and p_l represent here, and the valve in a feedback pneumatic circuit works as a pneumatic oscillator that generates periodic motion from an air source of constant pressure. It is a precedent for the hardware idea. It does not show that the first-order cycle model above describes their circuit, that their snap pressures play the role of p_h and p_l beyond the analogy, or that the coupling law derived in Chapter 29 has been realized.

Equation (2.1) shows the quantities available for tuning: the restriction sets τ_c, the vent sets τ_o, and the valve thresholds set the logarithmic factors. Because the period depends on the supply pressure p_s through the first logarithm, a change in p_s shifts the frequency. This gives a natural route by which a common supply line can act as a weak global drive, a point used in Part V.

2.2 From Chamber to Phase

Assign phase $\theta = 0$ to the moment the valve closes. Phase then advances through the charging interval and the discharge interval at rates that differ, so the oscillator is not uniform in the original time variable. A change of phase variable to a uniformly advancing one is always possible for a single oscillator, and this is the sense in which a pneumatic chamber enters the phase equations of Part II with natural frequency $\omega = 2\pi/T$.

Chapter 3

Pneumatic Elements as Couplers

A single chamber is an oscillator. Two chambers connected by a tube, a shared supply, or a mechanical link are coupled oscillators. The coupling strength is set by the conductance of the connection and the sign of the interaction by whether a rise in one chamber promotes or delays the valve closing in the other. These identifications are made precise in Part V. At this stage it suffices to note that a coupling in this setting is a physical object that can be adjusted, closed off or opened, which is what makes release a mechanical possibility and not only a mathematical one.

Part II

Coupling and Hierarchy

Chapter 4

Coupled Phase Oscillators

Scope of this part. The Shift Lemma and the two-sided coherence bounds are identities and elementary lemmas. The locking results are theorems about the phase model, and the linearized group estimate is a statement about a linearization. Nothing here concerns a chamber until the reduction to the phase model is bounded.

4.1 The Basic Model

Take N oscillators with phases θ_i, natural frequencies ω_i, coupling strengths $K_{ij} \geq 0$, preferred phase offsets α_{ij}, and external drives $I_i(t)$:

$$\dot{\theta}_i = \omega_i + \sum_j K_{ij} \sin\big(\theta_j - \theta_i - \alpha_{ij}\big) + I_i(t). \tag{4.1}$$

With $\alpha_{ij} = 0$ and $I_i = 0$ this is the Kuramoto model [7, 19]. Synchronization in networks of such phase oscillators, including phase cohesiveness, conditions that depend on the interconnection graph, and performance estimates, is surveyed by Dörfler and Bullo [2], and the cohesion results of Part III are of the type surveyed there. The offsets α_{ij} follow the phase-lag extension of Sakaguchi and Kuramoto [17]. The reason for carrying both K_{ij} and α_{ij} is that they do different jobs. The matrix K says how strongly two subsystems constrain one another. The matrix α says what relationship the constraint enforces. A gait is specified by a set of offsets, and the strength of the coupling says how firmly the gait is held.

4.2 Prescribed Patterns and the Shift Lemma

A coordination pattern for a group G is a set of target offsets β_i, so that the pattern is met when $\theta_i - \beta_i$ takes a common value for all $i \in G$. Offsets are admissible as pairwise preferences α_{ij} exactly when they come from such a β, that is, when $\alpha_{ij} = \beta_j - \beta_i$.

Lemma 4.1 (Shift)**.** *If $\alpha_{ij} = \beta_j - \beta_i$ and $\vartheta_i = \theta_i - \beta_i$, then Equation* (4.1) *becomes*

$$\dot{\vartheta}_i = \omega_i + \sum_j K_{ij} \sin(\vartheta_j - \vartheta_i) + I_i(t), \tag{4.2}$$

which is the plain Kuramoto model in the shifted variables.

Proof. $\sin(\theta_j - \theta_i - \alpha_{ij}) = \sin(\vartheta_j - \vartheta_i + \beta_j - \beta_i - \alpha_{ij}) = \sin(\vartheta_j - \vartheta_i)$. □

The lemma reduces every prescribed pattern, including antiphase and traveling-wave patterns, to the in-phase problem, and so transfers the classical synchronization results [7, 19] to gaits.

4.3 The Pattern-Adjusted Order Parameter

For a group G_a with target offsets β_i define

$$Z_a = r_a e^{\mathrm{i}\Psi_a} = \frac{1}{|G_a|} \sum_{i \in G_a} e^{\mathrm{i}(\theta_i - \beta_i)}. \tag{4.3}$$

Subtracting the target offsets matters. The unadjusted order parameter vanishes for an orderly antiphase pattern, and would report that pattern as incoherence.

Lemma 4.2 (Adequacy of r_a)**.** *$r_a = 1$ if and only if $\theta_i - \beta_i$ is the same for all $i \in G_a$, that is, if and only if the prescribed pattern holds exactly. If every $\theta_i - \beta_i$ lies within $\epsilon \le \pi/2$ of some reference angle c on the circle, for instance Ψ_a or the mean $\bar{\vartheta}_a$ of the lifted phases, then $r_a \ge \cos\epsilon$.*

Proof. r_a is the modulus of an average of unit vectors, so $r_a \le 1$, with equality exactly when all the vectors coincide. For the bound, $r_a = |Z_a| \ge \mathrm{Re}\big(e^{-\mathrm{i}c} Z_a\big) = \frac{1}{|G_a|} \sum_i \cos(\theta_i - \beta_i - c) \ge \cos\epsilon$ because cosine is decreasing on $[0, \pi]$ and $\epsilon \le \pi/2$ keeps every term nonnegative; the reference angle is arbitrary, which is why the lemma can be used with the mean phase. □

Two collective phases. The argument Ψ_a of the order parameter and the arithmetic mean $\bar{\vartheta}_a = |G_a|^{-1} \sum_{i \in G_a} \vartheta_i$ of the lifted pattern-adjusted phases are different quantities and are written with different symbols throughout. They are not equal in general: for the three phases $0, 0, \pi/2$ one has $\Psi = \arg(2 + \mathrm{i}) \approx 0.464$ and $\bar{\vartheta} = \pi/6 \approx 0.524$. The reduction of Part III is an identity for $\bar{\vartheta}_a$, because averaging the oscillator equations over a group produces the mean; Ψ_a enters only through r_a and the coherence results of this chapter. When all phases of a group lie in an arc of width $D < \pi$, both Ψ_a and $\bar{\vartheta}_a$ lie in that arc, so $|\Psi_a - \bar{\vartheta}_a| \le D$, provided the lifts are compatible, that is, chosen so that all lifted phases lie in one interval of length $D < \pi$ (a common arc of width below π also makes $Z_a \ne 0$, so that Ψ_a is defined; the argument of a sum of unit vectors in such an arc lies in the arc), and any statement transferred from one to the other carries that difference.

Lemma 4.3 (Spread and coherence)**.** *Let n pattern-adjusted phases have order parameter $re^{\mathrm{i}\Psi}$. (a) If the phases lie in an arc of width $D \le \pi$ then $r \ge \cos(D/2)$, so $D \ge 2\arccos r$. (b) If $n(1-r) < 1$ then every phase lies within circular distance $\arccos\big(1 - n(1-r)\big) < \pi/2$ of Ψ, so the phases lie in an arc of width at most $2\arccos\big(1 - n(1-r)\big) < \pi$.*

Proof. (a) Let c be the center of the arc, so every phase is within $D/2$ of c. Then $nr \ge \big|\sum_j e^{\mathrm{i}(\theta_j - c)}\big| \ge \sum_j \cos(\theta_j - c) \ge n\cos(D/2)$. (b) Since the imaginary part of $\sum_j e^{\mathrm{i}(\theta_j - \Psi)}$ vanishes, $\sum_j \cos(\theta_j - \Psi) = nr$, hence $\sum_j \big(1 - \cos(\theta_j - \Psi)\big) = n(1-r)$. Every term is nonnegative, so each is at most $n(1-r)$, which gives $\cos(\theta_i - \Psi) \ge 1 - n(1-r)$ for every i. □

The two bounds do not meet, and no bound free of n exists: $n-1$ phases at 0 and one at π have $r = 1 - 2/n$ and spread π. A coherence close to one therefore bounds the spread only through the group size.

The lemma establishes what r_a detects: it detects the pattern, and it bounds the worst error from below as a function of r_a only in one direction. A small r_a implies a large error somewhere; a large r_a does not identify which members carry the residual error. Chapter 17 returns to what this cannot distinguish.

4.4 Two Oscillators

Let two oscillators be coupled symmetrically with strength K and preferred offset α. Write $\varphi = \theta_2 - \theta_1$ and $\Delta\omega = \omega_2 - \omega_1$. Then

$$\dot{\varphi} = \Delta\omega - 2K\sin(\varphi - \alpha). \tag{4.4}$$

Proposition 4.1. *Equation* (4.4) *has phase-locked solutions if and only if* $|\Delta\omega| \le 2K$. *For* $|\Delta\omega| < 2K$ *the locked offset* $\varphi^* = \alpha + \arcsin\big(\Delta\omega/2K\big)$ *is exponentially stable, with linearized rate* $2K\cos(\varphi^* - \alpha) > 0$. *At* $|\Delta\omega| = 2K$ *the two branches merge into a single locked offset* $\alpha \pm \pi/2$, *the linearized restoring rate vanishes, and the offset is semistable (attracting from one side only), so it is not stable in the strict sense. At* $\Delta\omega = 0$ *the offset equals* α *exactly.*

Proof. A fixed point requires $\sin(\varphi - \alpha) = \Delta\omega/2K$, which has solutions exactly when $|\Delta\omega| \le 2K$. Linearizing, $\dot{\delta} = -2K\cos(\varphi^* - \alpha)\,\delta$, which is exponentially stable when $\cos(\varphi^* - \alpha) > 0$, selecting the arcsine branch, and this holds exactly when $|\Delta\omega| < 2K$. At $|\Delta\omega| = 2K$ the cosine vanishes; the second-order expansion $\dot{\delta} = K\,\mathrm{sgn}(\Delta\omega)\,\delta^2 + O(\delta^3)$ shows that perturbations to one side return and to the other side escape, so the fixed point is semistable. □

Locking is a threshold phenomenon in the ratio of frequency mismatch to coupling, so reducing K eventually releases the pair into drift, and the locked offset is set by α up to a correction proportional to the mismatch.

4.5 Coordination Within a Group

Take a group of n oscillators with all-to-all coupling κ and write $\bar{\omega}$ for the mean frequency and $d_i = \omega_i - \bar{\omega}$. In shifted variables the group obeys $\dot{\vartheta}_i = \omega_i + \kappa n\, r \sin(\Psi - \vartheta_i)$, where $re^{\mathrm{i}\Psi}$ is the adjusted order parameter (4.3).

Proposition 4.2 (Identical frequencies)**.** *If all* $d_i = 0$, *then on the set* $r > 0$, *where* Ψ *is defined,* $\dot{r} = \kappa r \sum_i \sin^2(\vartheta_i - \Psi) \ge 0$, *so* r *is nondecreasing along trajectories there. The incoherent set* $r = 0$ *consists of rest points, since every* $\dot{\vartheta}_i$ *then vanishes, so* $r = 0$ *is invariant and the result says nothing about trajectories that start on it.*

Proof. Since the sum of $\sin(\vartheta_i - \Psi)$ over the group is zero, $\dot{\Psi}$ drops out of $\dot{r} = -\frac{1}{n}\sum_i \sin(\vartheta_i - \Psi)(\dot{\vartheta}_i - \dot{\Psi})$, leaving $\dot{r} = \frac{1}{n}\sum_i \sin(\vartheta_i - \Psi)\,\kappa n r \sin(\vartheta_i - \Psi)$. □

Convergence of almost every initial condition to $r = 1$ in this identical case is the classical result [19]. For the nonidentical case the following is established only for the linearization.

Proposition 4.3 (Linearized group estimate)**.** *Let $e_i = \vartheta_i - \bar{\vartheta}$ be the deviation of member i from the group mean phase. In the small-angle regime where $\sin x \approx x$ the deviations obey $\dot{e}_i = d_i - \kappa n\, e_i$, so that*

$$|e_i(t)| \le e^{-\kappa n t}|e_i(0)| + \frac{|d_i|}{\kappa n}\big(1 - e^{-\kappa n t}\big). \tag{4.5}$$

The estimate has the form wanted for the coordination theorem, an exponential transient plus a term proportional to the disturbance, here frequency mismatch. It is a statement about the linearized model. The nonlinear statement is Theorem 6.1 in Part III.

Chapter 5

Block-Structured Coupling

5.1 Hierarchy as Structure of K

A hierarchy is encoded by partitioning the oscillators into functional groups, such as the joints of a finger, the fingers of a hand, and the hand and arm of a limb. Let $g(i)$ denote the group of oscillator i. Define

$$K_{ij} = \begin{cases} \kappa_{\text{in}} & g(i) = g(j), \\ \kappa_{\text{out}}(g(i), g(j)) & g(i) \neq g(j), \end{cases} \qquad \kappa_{\text{out}} \ll \kappa_{\text{in}}. \tag{5.1}$$

Strong coupling inside groups and weak coupling between them is the structure used for nested motor oscillation in the author's earlier writing on gesture and locomotion [26, 25]. Recursion gives a hierarchy of more than two levels: groups of groups, with coupling weaker at each coarser level.

5.2 Why the Separation of Scales Matters

When κ_{out} is small relative to κ_{in}, each group synchronizes internally on a fast timescale, and the groups then interact through their collective phases $\bar{\vartheta}_g$ on a slow timescale. This motivates describing the system at two levels, and Theorems 6.1 and 7.1 give the quantitative version under their stated hypotheses: a fast description in which each group is a single effective oscillator with phase $\bar{\vartheta}_g$, and a slow description of the coupling among those effective oscillators. The condition is a statement about the ratio of timescales and must be checked in any particular body. It is the mathematical form of the claim that a limb is a unit.

Remark 5.1. Central pattern generators with this kind of modular organization are well documented in animals and have been used as controllers for robots [11, 5, 6]. Symmetry-based accounts of gaits [4] supply a complementary route to the offsets α_{ij}, since the symmetry group of the network constrains which offset patterns can occur.

Part III

Synchronization and Release

Chapter 6

Maintained Coordination

Scope of this part. The cohesion, reduction, capture, principal, adaptation and three-group results are theorems about the phase model, stated as sufficient conditions with their excluded regions, and several are conservative. Their hypotheses are: symmetric internal weights or lower-bounded weights on complete graphs; uniform cross coupling in the uniform case and reciprocal cross weights in the nonuniform case; consistent targets; a constant or bounded frequency mismatch; and initial conditions inside the invariant set where an invariance statement is made. The transition hypotheses and the slow-adaptation law are proposals, and the boundary simulations are numerical findings at the stated parameters.

Part III treats three operations on a coupled hierarchy and keeps them apart. Maintenance holds a set of phase relationships. Release relaxes some relationships while preserving others. Transition replaces one set of relationships by another. The part ends with a single theorem, proved for two groups with uniform coupling, that assembles maintenance, release and recapture into one statement with explicit error bounds.

Throughout, a group is a set of oscillators with a prescribed pattern β_i, and the working variables are the lifted pattern-adjusted phases $\vartheta_i = \theta_i - \beta_i \in \mathbb{R}$ of the Shift Lemma. The spread of a group is $D_a(t) = \max_{i\in a} \vartheta_i - \min_{i\in a} \vartheta_i$. Small spread means the group keeps its pattern. The spread is a property of the lift, and hypotheses on it, such as $D(0) < \pi - \gamma$, are statements about a choice of representatives in $\mathbb{R}$. A circular configuration satisfies such a hypothesis exactly when its pattern-adjusted phases lie in an arc of the stated width, and the lift is then chosen inside that arc; a configuration spread around the circle admits no such lift and is outside the theorem. A bounded spread is cohesion of the group and is not by itself phase locking: locking means that the relative phases converge to constants, which the cohesion theorems do not assert and which is treated separately where it is claimed.

6.1 Cohesion of a Group

Consider a group of n oscillators with uniform internal coupling κ,

$$\dot{\vartheta}_i = \omega_i + I_i(t) + \kappa \sum_{j\in g} \sin(\vartheta_j - \vartheta_i) + \varepsilon_i(t), \tag{6.1}$$

where $\varepsilon_i(t)$ is whatever the rest of the body does to oscillator i, bounded by $|\varepsilon_i(t)| \le E$. Let W be the supremum over time of the spread of the effective frequencies $\omega_i + I_i(t)$.

Theorem 6.1 (Cohesion). *Suppose $\kappa n > W + 2E$ and set $\gamma = \arcsin\big((W+2E)/(\kappa n)\big) \in (0, \pi/2)$. (a) If $D(0) \le \gamma$ then $D(t) \le \gamma$ for all $t \ge 0$. (b) If $D(0) < \pi - \gamma$, then D is nonincreasing wherever $\gamma \le D \le \pi - \gamma$, $\limsup_{t\to\infty} D(t) \le \gamma$, and for every $\eta > 0$ with $\gamma + \eta < D(0)$ there is a finite time $T(D(0), \eta)$ after which $D \le \gamma + \eta$ for ever. The time T is not bounded uniformly over $D(0) < \pi - \gamma$.*

Proof. Let $M(t)$ and $m(t)$ be the largest and smallest of the $\vartheta_i(t)$. The upper right Dini derivative of $D = M - m$ is at most $\dot\vartheta_p - \dot\vartheta_q$ for indices p, q attaining the maximum and minimum. The effective frequencies of p and q, including drive and external terms, differ by at most $W + 2E$, and the coupling terms give

$$\kappa \sum_j \big[\sin(\vartheta_j - M) - \sin(\vartheta_j - m)\big] = -\kappa \sum_j \big[\sin a_j + \sin b_j\big],$$

with $a_j = M - \vartheta_j \ge 0$, $b_j = \vartheta_j - m \ge 0$ and $a_j + b_j = D$. For $D \le \pi$, $\sin a_j + \sin b_j = 2\sin(D/2)\cos\big((a_j - b_j)/2\big) \ge 2\sin(D/2)\cos(D/2) = \sin D$, because $|a_j - b_j| \le D$ and $D/2 \le \pi/2$. Hence

$$D^+ D \le (W + 2E) - \kappa n \sin D \qquad \text{whenever } D \le \pi. \tag{6.2}$$

At $D = \gamma$ the right side is zero, and on $[\gamma, \pi - \gamma]$ it is at most zero because $\sin D \ge \sin\gamma$ there. This gives (a), and shows that $\{D \le \gamma'\}$ is invariant for every $\gamma' \in [\gamma, \pi - \gamma]$ and that D is nonincreasing on $[\gamma, \pi - \gamma]$. For (b), fix $\eta > 0$ with $\gamma + \eta < D(0)$. On $[\gamma + \eta, D(0)]$ concavity of the sine and $D(0) < \pi - \gamma$ give $\sin D \ge s_\eta := \min\{\sin(\gamma + \eta), \sin D(0)\} > \sin\gamma$. Hence $D^+ D \le -c$ with $c = \kappa n(s_\eta - \sin\gamma) > 0$, and the interval is crossed in time at most $(D(0) - \gamma - \eta)/c$. As $D(0) \to \pi - \gamma$ the constant c tends to zero, so no bound uniform over the whole set can be obtained from this argument. After the crossing, invariance of $\{D \le \gamma + \eta\}$ gives the remaining claims. □

The theorem contains two different kinds of statement and they should not be confused. The invariance statement says that a trajectory that starts inside the set $\{D \le \gamma\}$ never leaves it. It depends on an assumption about the initial condition and says nothing about trajectories that start outside. The attraction statement says that a trajectory starting anywhere in $\{D < \pi - \gamma\}$ enters every set $\{D \le \gamma + \eta\}$ in finite time and stays in it, which is attraction to the set $\{D \le \gamma\}$ and not arrival in it; it requires no initial assumption inside that larger set, but it does not say how long the approach takes beyond the finite bound from the proof.

The bound is the nonlinear form of the linearized estimate of Part II: for small $(W + 2E)/\kappa n$, $\gamma \approx (W + 2E)/\kappa n$. It states the separation required among internal coupling, external coupling and frequency dispersion.

Remark 6.1 (What lowering external coupling does). With W, κ and n fixed, γ is nondecreasing in E. Lowering the external coupling therefore cannot worsen the worst-case guarantee. This is a statement about the bound and not about the trajectory. External coupling can help coordination: take two members with frequency mismatch 0.5 and internal coupling $\kappa = 0.3$, whose locked spread is $\arcsin(0.5/0.6) \approx 0.99$, and let the external influence act as a drive that reduces the mismatch to 0.3; the locked spread falls to $\arcsin(0.3/0.6) \approx 0.52$. The theorem treats external influence as an adversarial disturbance of size E and so cannot see this, and the boundary simulations of Chapter 16 show the effect: a group joined to a

cohesive neighbor stayed cohesive at internal couplings where the same group in isolation lost cohesion.

Corollary 6.1 (Sensory correction)**.** *Suppose the term $\varepsilon_i(t)$ of Equation* (6.1) *contains, besides the external coupling bounded by E, a sensory contribution $Z_i(\theta_i)\xi_i(t)$ with $|Z_i(\theta_i)\xi_i(t)| \le F$, where ξ_i is a sensory perturbation and Z_i its phase sensitivity. Then Theorem 6.1 holds with E replaced by $E + F$: the condition becomes $\kappa n > W + 2E + 2F$ and*

$$\gamma = \arcsin\!\left(\frac{W + 2E + 2F}{\kappa n}\right),$$

with the same initial-condition assumptions. In the graph version, Φ_a gains the Euclidean norm of the vector of sensory bounds.

Proof. The sensory term is one more bounded contribution to ε_i, so $|\varepsilon_i| \le E + F$, and the theorem uses ε_i only through this bound. □

Remark 6.2. The form $\dot\theta_i = \omega_i + \sum_j H_{ij}(\theta_j - \theta_i) + Z_i(\theta_i)\xi_i(t)$ is a modeling translation of the physiological account in the review of Duysens and Van de Crommert [3], in which sensory input can act on the generator to aid transitions between phases of the step cycle and reflex effects depend on the stage of the cycle. The equation is not established by that review, and the corollary does not claim that the bound explains biological locomotion. The bound treats the sensory term as an adversarial disturbance, whereas the review describes sensory input that helps the rhythm, so the guarantee is conservative in the same sense as the remark above. The body path of Chapter 24, a pressure perturbation weighted by the phase response, has the same form with $F = \zeta_b$.

Chapter 7

The Effective Oscillator

7.1 One-Level Reduction

Let the body consist of groups $a = 1, \dots, m$ of sizes n_a, with uniform internal coupling κ_a and uniform cross coupling $\kappa_{ab} = \kappa_{ba} \geq 0$ between groups a and b. Let the preferred offsets be $\alpha_{ij} = \beta_j - \beta_i + a_{ab}$ for $i \in a$, $j \in b$, with group offsets $a_{ba} = -a_{ab}$. Define the collective phase of group a as the mean of its lifted phases, $\bar{\vartheta}_a = n_a^{-1} \sum_{i \in a} \vartheta_i$, and write $\bar{\omega}_a$ for the mean of $\omega_i + I_i$ over the group.

Theorem 7.1 (Reduction)**.** *The collective phases satisfy exactly*

$$\dot{\bar{\vartheta}}_a = \bar{\omega}_a + \sum_{b \neq a} n_b \kappa_{ab} \sin(\bar{\vartheta}_b - \bar{\vartheta}_a - a_{ab}) + \rho_a(t), \tag{7.1}$$

with

$$|\rho_a(t)| \leq \sum_{b \neq a} n_b \kappa_{ab} \big(D_a(t) + D_b(t) \big). \tag{7.2}$$

Proof. Average the equation for ϑ_i over $i \in a$. The internal coupling contributes

$$\kappa_a n_a^{-1} \sum_{i,j \in a} \sin(\vartheta_j - \vartheta_i) = 0$$

by antisymmetry. The cross terms are

$$n_a^{-1} \sum_{i \in a} \sum_{b} \sum_{j \in b} \kappa_{ab} \sin(\vartheta_j - \vartheta_i - a_{ab}).$$

Replace ϑ_j by $\bar{\vartheta}_b$ and ϑ_i by $\bar{\vartheta}_a$; the argument changes by at most $D_a + D_b$, and $|\sin x - \sin y| \leq |x - y|$, which gives the remainder and the bound. □

The reduced equation has the form wanted for a higher level: each group is an effective oscillator with its own frequency and sinusoidal coupling to the others, and the error is explicit and proportional to the spreads. When Theorem 6.1 holds, $D_a \leq \gamma_a$ and the error bound is a constant. A hierarchy with more than two levels applies the theorem repeatedly, each level carrying the remainder of the one below, so that the accumulated error is a finite sum of such terms. The bookkeeping for more than one nesting is not carried out here.

Chapter 8

Release

Definition 8.1. Release of a group g at a given level of the hierarchy is a change in the coupling structure or drive such that the cross coherence

$$c_{gh} = \left| \left\langle e^{\mathrm{i}(\bar{\vartheta}_h - \bar{\vartheta}_g)} \right\rangle_{\Delta t} \right| \tag{8.1}$$

falls for the groups h outside g, averaged over a window long compared with the slow period, while the spread D_g stays small.

The definition is relative to a level: an arm becomes independent of the gait while its fingers, wrist and elbow remain coordinated with one another. It describes what happens to two measurements, not how it is brought about.

8.1 Routes to Release

For two groups write $N = (n_g + n_h)\kappa_{gh}$ and $\Delta\bar{\omega} = \bar{\omega}_h - \bar{\omega}_g$. Neglecting ρ, Equation (7.1) gives, for $u = \bar{\vartheta}_h - \bar{\vartheta}_g - a_{gh}$,

$$\dot{u} = \Delta\bar{\omega} - N \sin u, \tag{8.2}$$

which is Equation (4.4) with K replaced by $N/2$ and locks exactly when $|\Delta\bar{\omega}| \leq N$.

Two routes release the pair, decoupling and detuning, and a third operation, changing the offset, does not. Release by decoupling lowers κ_{gh} so that $N < |\Delta\bar{\omega}|$ with the internal couplings unchanged. Release by detuning leaves the couplings alone and changes $\Delta\bar{\omega}$: a constant drive $I_i = \delta$ on every member of group g adds δ to its mean frequency and so can push the pair outside the locking region. Changing the offset a_{gh} does not appear in the locking condition and therefore is not a route to release; it changes which relationship is held and is the operation treated as transition in Chapter 10.

Hypothesis 8.1. *In a given body, which of decoupling and detuning accounts for observed release is an empirical question. Under decoupling a measured conductance or link stiffness between the groups changes at release, and under detuning it does not.*

Chapter 9

Independent Action and Recapture

A released group follows its own drive. Its members remain locked to one another by Theorem 6.1 and its collective phase advances at its own rate. Nothing requires the released group to be idle; independent action is motion governed by the group's own structure and not by the larger rhythm.

9.1 The Capture Theorem

Recapture is the reverse of release. With the error term of Theorem 7.1 included, the relative phase obeys $\dot{u} = \Delta\bar{\omega} - N\sin u + r(t)$ with $|r| \le R$.

Theorem 9.1 (Capture)**.** *Let $s^* = (|\Delta\bar{\omega}| + R)/N < 1$ and $\gamma_* = \arcsin s^*$. (a) For every $\gamma' \in [\gamma_*, \pi - \gamma_*)$ the arc $[-\gamma', \gamma']$ of the circle is forward invariant. (b) If $|u(0)| < \pi - \gamma_*$ (taken in $(-\pi, \pi]$), then $\limsup_{t\to\infty} |u(t)| \le \gamma_*$, and for every $\eta > 0$ with $\gamma_* + \eta < |u(0)|$ the phase enters $[-\gamma_* - \eta, \gamma_* + \eta]$ within time $(|u(0)| - \gamma_* - \eta)/c$, where $c = N\big[\min\{\sin(\gamma_* + \eta), \sin|u(0)|\} - s^*\big] > 0$, and remains there. (c) If $|u(0)| \le \pi - \gamma_* - \eta$ this time is at most π/c_η with $c_\eta = N\big[\sin(\gamma_* + \eta) - s^*\big]$, uniformly in $u(0)$. No finite-time entry into the exact arc $[-\gamma_*, \gamma_*]$ is claimed. The allowable frequency mismatch is $|\Delta\bar{\omega}| < N - R$.*

Proof. Where $\sin u > s^*$, $\dot{u} \le |\Delta\bar{\omega}| + R - N\sin u < 0$, and where $\sin u < -s^*$, $\dot{u} > 0$. For $\gamma' \in [\gamma_*, \pi - \gamma_*)$ we have $\sin\gamma' \ge s^*$, so $\dot{u} \le 0$ at $u = \gamma'$ and $\dot{u} \ge 0$ at $u = -\gamma'$, which gives (a). For (b), take $u(0) \in [\gamma_* + \eta, \pi - \gamma_*)$ (the negative case is symmetric). On $[\gamma_* + \eta, u(0)]$ concavity of the sine gives $\sin u \ge \min\{\sin(\gamma_* + \eta), \sin u(0)\}$, so $\dot{u} \le -c$ and the interval is crossed within time $(u(0) - \gamma_* - \eta)/c$; invariance then keeps u in the enlarged arc, and since η is arbitrary $\limsup|u| \le \gamma_*$. For (c), when $u(0) \le \pi - \gamma_* - \eta$ concavity gives $\sin u \ge \sin(\gamma_* + \eta)$ on the whole interval, so $c \ge c_\eta$ and the crossing time is at most π/c_η. The simplest case $\dot{u} = -N\sin u$, with $\Delta\bar{\omega} = R = 0$, shows that no uniform bound is possible near $|u(0)| = \pi$: the time to reach η from u_0 is $N^{-1}\ln\big(\tan(u_0/2)/\tan(\eta/2)\big)$, which diverges as $u_0 \to \pi$, and the exact arc is the single point $\{0\}$, approached only asymptotically. □

The capture region is the whole circle except a band of half-width γ_* around the antipode $u = \pi$. A relative phase that starts inside that band is not covered by this argument. In the remainder-free equation $\dot{u} = \Delta\bar{\omega} - N\sin u$ with $|\Delta\bar{\omega}| < N$ there are two equilibria, a stable one at $a = \arcsin(\Delta\bar{\omega}/N)$ and an unstable one at $\pi - a$ modulo 2π, and the unstable one

lies in the band; a trajectory started exactly there is not captured, and one started near it may escape in either direction, with a single phase slip in the process. For a released group that has drifted during release, whether its relative phase lands in the band is a question of timing, and a practical recapture schedule either waits for u to leave the band or applies a brief detuning.

Two phenomena of gesture appear here [26]. Anticipation is the early onset of coupling to the next pattern while the current one is still being released. Carry-over is the persistence of the previous relationship in the released group, which near threshold is slow: when $|\Delta\bar{\omega}|$ exceeds N by only a small amount, u passes slowly through the region where $\sin u$ is near $\Delta\bar{\omega}/N$, and the passage time grows without bound as the margin shrinks.

Chapter 10

Transition Between Coordination Patterns

10.1 Four Candidate Mechanisms

A transition moves the system from one stable coordination pattern to another. It is tempting to assume that a transition passes through generalized desynchronization. That is one hypothesis among several, and this chapter keeps them separate so that observations can tell them apart.

Hypothesis 10.1 (Generalized desynchronization)**.** *The transition proceeds by weakening selected entries of K, allowing reorganization, and restoring coupling with a new α.*

Hypothesis 10.2 (Continuous offset change)**.** *The preferred offset changes continuously and the locked offset follows it without loss of locking.*

Hypothesis 10.3 (Phase slip)**.** *During the transition one oscillator gains or loses a full cycle relative to another, so that the phase difference changes by 2π in addition to the change in target.*

Hypothesis 10.4 (Attractor switch)**.** *For the same K and α the network has several stable offset patterns, and a transient drive moves the state across a basin boundary from one to another.*

Remark 10.1 (Other routes to a change of gait)*.* Release and recapture is one way to change a movement and not the only established one. Controllers for fish robots change gait through command-dependent frequency and amplitude settings and through the choice of which propulsors contribute, with two input parameters regulating speed, gait and maneuvering [22]. In the terms of this monograph these are changes of the natural frequencies, of the amplitudes and engagement of Chapter 11, and of the drive, and none requires weakening a coupling. The four hypotheses below concern changes of the preferred pattern within a coupled hierarchy and do not claim that gait changes in general pass through release.

10.2 Conditions

Theorem 9.1 already covers part of this. If the group offset a_{gh} is changed while coupling is kept above the capture threshold, then u jumps by the change in a_{gh} and relaxes toward $[-\gamma_*, \gamma_*]$, entering any enlarged arc in finite time, provided the new u lies outside the antipodal band. This is a transition that preserves locking in the sense of recapture without release. A large change that places u in the band, or one made while coupling is below threshold, requires either temporary loss of lock or a slip. Continuous offset change keeps the system locked when $a_{gh}(t)$ is of class C^1 and moves slowly enough that the equation $\dot{u} = \Delta\bar{\omega} - \dot{a}_{gh} - N \sin u + r$ keeps $(\sup |\Delta\bar{\omega} - \dot{a}_{gh}| + R)/N < 1$, which is Theorem 9.1 with the effective mismatch $\Delta\bar{\omega} - \dot{a}_{gh}$ (the moving-target statement of Chapter 14); this is a sufficient condition on the speed and not a statement of the tracking error. The second and third hypotheses are therefore related by the speed of the change, and an experiment must vary the speed.

10.3 Discriminating Observations

The mechanisms predict different time courses of c_{gh} and $\Phi = \bar{\vartheta}_h - \bar{\vartheta}_g$. Under generalized desynchronization c_{gh} dips toward zero during the transition and recovers. Under continuous offset change c_{gh} stays high and Φ moves smoothly. Under phase slip c_{gh} dips briefly and Φ shows a single 2π excursion. Under attractor switch c_{gh} is high before and after, Φ jumps between two discrete values, and the jump is preceded by an identifiable drive pulse. No claim about which one occurs in a gait change is made here in advance of such measurements.

Chapter 11

Stopping, Amplitude and Engagement

Loss of synchrony does not stop movement; it can leave several subsystems moving independently, which is release. Conversely, a body can cease visible movement while an internal timing process continues. The theory therefore separates three variables that the phase equations alone conflate.

Phase θ_i is the timing variable. Amplitude $a_i \geq 0$ scales the excursion of the oscillator. Engagement $e_i \in [0, 1]$ is the fraction of the oscillator's output transmitted to the mechanism. A simple output model is

$$y_i(t) = e_i(t)\, a_i(t)\, s\big(\theta_i(t)\big), \tag{11.1}$$

where s is a 2π-periodic waveform. Movement stops when y_i vanishes, which can happen because $a_i = 0$, because $e_i = 0$, or because the oscillation itself has ceased. Only the last is a change in the phase dynamics; the first two leave θ_i advancing and so leave coordination intact. The distinction between timing, amplitude and which units contribute to the output has a precedent in fish-robot controllers, where gait changes are made by amplitude and frequency settings and by selecting which fins act [22]. Stopping with $e_i = 0$ keeps the system ready to resume in step: by Theorem 9.1 with engagement restored the pattern is approached asymptotically, provided the phase relations were never released and the relative phase stays outside the excluded band. This is the sufficient condition for stopping and restarting in the model with engagement variables; it does not describe the mechanical transients of re-engagement.

In the pneumatic implementation these variables correspond to different acts. Venting the bladder or slackening a cord changes engagement. Reducing supply pressure lowers amplitude, but by Equation (2.1) it also changes the period, so supply pressure couples amplitude to frequency and cannot be treated as a pure amplitude control. Closing off the cycling valve stops the oscillation itself.

Chapter 12

The Principal Theorem

The results of this part combine into a single statement for two groups with uniform coupling.

Theorem 12.1 (Hierarchical release and recapture)**.** *Let groups g and h of sizes n_g, n_h have uniform internal couplings κ_g, κ_h, consistent target offsets, group offset a_{gh}, and uniform external coupling $\kappa_{gh}(t) \in [0, \bar{\kappa}]$, a time-varying control. Let W_a be the spread of effective frequencies within group a, and $\Delta\bar{\omega}$ the constant difference of group mean frequencies. Define $E_g = n_h\bar{\kappa}$, $E_h = n_g\bar{\kappa}$ and assume $\kappa_a n_a > W_a + 2E_a$ for $a = g, h$. Set $\gamma_a = \arcsin\big((W_a + 2E_a)/(\kappa_a n_a)\big)$, $\bar{N} = (n_g + n_h)\bar{\kappa}$, and assume $s^* = |\Delta\bar{\omega}|/\bar{N} + \gamma_g + \gamma_h < 1$ with $\gamma_* = \arcsin s^*$. Suppose $D_a(0) \le \gamma_a$. Then: (i) Internal coordination is preserved: $D_a(t) \le \gamma_a$ for all $t \ge 0$, whatever the schedule $\kappa_{gh}(t)$. (ii) Release: on any interval where κ_{gh} is held at a value κ_{rel} with $N_{\mathrm{rel}}(1+\gamma_g+\gamma_h) < |\Delta\bar{\omega}|$, where $N_{\mathrm{rel}} = (n_g+n_h)\kappa_{\mathrm{rel}}$, the relative phase u drifts monotonically at rate at least $|\Delta\bar{\omega}| - N_{\mathrm{rel}}(1 + \gamma_g + \gamma_h)$, so the groups are not locked. (iii) Recapture: if κ_{gh} is restored to $\bar{\kappa}$ at a time t_r at which $|u(t_r)| < \pi - \gamma_*$, then $\limsup_{t\to\infty} |u(t)| \le \gamma_*$; moreover for every $\eta > 0$ with $\gamma_* + \eta < |u(t_r)|$ the phase enters $[-\gamma_* - \eta, \gamma_* + \eta]$ within a finite time given in Theorem 9.1, uniformly bounded by π/c_η when $|u(t_r)| \le \pi - \gamma_* - \eta$, and remains there. Consequently every pair $i \in g$, $j \in h$ satisfies $\limsup_{t\to\infty} |\vartheta_j - \vartheta_i - a_{gh}| \le \gamma_* + \gamma_g + \gamma_h$ modulo 2π, and for each $\eta > 0$ the bound $\gamma_* + \gamma_g + \gamma_h + \eta$ holds after a finite time.*

Proof. Part (i): the external perturbation on any member of group a is at most E_a in modulus, so Theorem 6.1 applies with those values and gives invariance of $D_a \le \gamma_a$; the bound uses the supremum $\bar{\kappa}$ and so holds for every schedule. For the rest, Theorem 7.1 with $D_a \le \gamma_a$ gives the equation for $u = \bar{\vartheta}_h - \bar{\vartheta}_g - a_{gh}$, $\dot{u} = \Delta\bar{\omega} - N(t)\sin u + r(t)$ with $N(t) = (n_g + n_h)\kappa_{gh}(t)$ and $|r| \le N(t)(\gamma_g + \gamma_h)$. Part (ii): $|\dot{u}|$ is bounded below by $|\Delta\bar{\omega}| - N_{\mathrm{rel}} - N_{\mathrm{rel}}(\gamma_g + \gamma_h) > 0$ with the sign of $\Delta\bar{\omega}$. Part (iii): with $N = \bar{N}$ and $R = \bar{N}(\gamma_g+\gamma_h)$ we have $(|\Delta\bar{\omega}|+R)/N = s^*$ and Theorem 9.1 applies. The final statement combines $|u| \le \gamma_*$ with the two within-group deviations from the means. □

The theorem establishes, for two groups with uniform coupling, that a selected subsystem can be released while its internal coordination remains within an explicit bound, and that the prescribed collective pattern is recovered with an explicit error, under stated margins on internal coupling, frequency spread and external coupling. The next three chapters remove the uniformity assumption for two groups, extend the argument to three levels, and test the boundaries of the resulting estimates by simulation. The phase model itself is a reduction of the underlying oscillators, and the theorems are statements about that model.

Chapter 13

Nonuniform Coupling

13.1 Setting

Let groups g and h have sizes n_g, n_h. Within group a the internal coupling is a symmetric nonnegative weight matrix $K^a = (K^a_{ij})$ with weighted Laplacian $L_a = \operatorname{diag}(K^a\mathbf{1}) - K^a$, and the internal graph is connected, so that its algebraic connectivity λ_a, the second smallest eigenvalue of L_a, is positive. The cross coupling is $K_{ij}(t) = \mu(t)\bar{K}_{ij}$ for $i \in g$, $j \in h$, with $\bar{K}_{ij} \geq 0$ and a control $\mu(t) \in [0,1]$ that scales all external coupling together. Write $\bar{S} = \sum_{i\in g, j\in h} \bar{K}_{ij}$, $\bar{N} = \bar{S}(1/n_g + 1/n_h)$, and let $\bar{d}^a_i$ be the total cross coupling of oscillator $i \in a$, so $\bar{d}^a$ is a vector with $\|\bar{d}^a\|_2$ its Euclidean norm. Let $f_i = \omega_i + I_i(t)$ be the effective frequencies and

$$\Omega_a = \sup_t \big\| f - \bar{f}_a \mathbf{1} \big\|_{2,a}, \qquad \Phi_a = \Omega_a + \|\bar{d}^a\|_2, \tag{13.1}$$

where $\bar{f}_a$ is the group mean and the norm is over the members of a. The quantity Ω_a measures frequency dispersion in the Euclidean norm, and $\|\bar{d}^a\|_2$ measures the worst-case external push. The cross coupling is reciprocal: member $j \in h$ receives weight $K_{ij}(t)$ from $i \in g$ exactly as i receives from j. The mean frequency mismatch $\Delta\bar{\omega} = \bar{f}_h - \bar{f}_g$ is assumed constant in the theorems below; the release statements need only a positive lower bound on $|\Delta\bar{\omega}|$ and the capture statements only an upper bound.

13.2 Cohesion Through the Internal Graph

Let $e_i = \vartheta_i - \bar{\vartheta}$ be the deviation of member i from the group mean of the lifted phases, and let D be the spread as before.

Theorem 13.1 (Graph cohesion)**.** *Suppose $\sqrt{2}\,\Phi_a < \lambda_a$ and set $\gamma_a = \arcsin\big(\sqrt{2}\,\Phi_a/\lambda_a\big)$. (a) If $\|e(0)\|_2 \leq \gamma_a/\sqrt{2}$ then $\|e(t)\|_2 \leq \gamma_a/\sqrt{2}$ for all t, and hence $D(t) \leq \gamma_a$. (b) If $\|e(0)\|_2 < (\pi-\gamma_a)/\sqrt{2}$, then $\|e(t)\|_2$ is nonincreasing wherever $\gamma_a/\sqrt{2} \leq \|e\|_2 \leq (\pi-\gamma_a)/\sqrt{2}$, $\limsup \|e(t)\|_2 \leq \gamma_a/\sqrt{2}$, and for every $\eta > 0$ with $\gamma_a + \eta < \sqrt{2}\|e(0)\|_2$ the norm enters and stays in the ball of radius $(\gamma_a+\eta)/\sqrt{2}$ after a finite time depending on $\|e(0)\|_2$ and η.*

Proof. Write $\sin(\vartheta_j - \vartheta_i) = c_{ij}(\vartheta_j - \vartheta_i)$ with $c_{ij} = \operatorname{sinc}(\vartheta_j - \vartheta_i)$, which is symmetric in i, j. Because the weights are symmetric, the coupling terms sum to zero over the group, so the group mean evolves by the mean of $f_i + \varepsilon_i$ alone, and the deviations obey $\dot{e} = \delta - L_c(t)\, e$,

where δ is the deviation of $f+\varepsilon$ from its mean and L_c is the Laplacian with weights $K^a_{ij}c_{ij}$. Since $\|\delta\|_2 \le \Omega_a + \|\varepsilon\|_2 \le \Phi_a$, we have for $V = \frac{1}{2}\|e\|_2^2$ and $r = \|e\|_2$ that $\dot V \le r\Phi_a - e^\top L_c e$. The spread satisfies $D \le \sqrt{2}\, r$, because $(e_i - e_j)^2 \le 2(e_i^2 + e_j^2) \le 2r^2$. For $\sqrt{2}r \le \pi$ the sinc function is decreasing, so $c_{ij} \ge \operatorname{sinc}(\sqrt{2}r)$ and $e^\top L_c e \ge \operatorname{sinc}(\sqrt{2}r)\, e^\top L_a e \ge \operatorname{sinc}(\sqrt{2}r)\lambda_a r^2$, since e has zero sum. Therefore $\dot r \le \Phi_a - (\lambda_a/\sqrt{2})\sin(\sqrt{2}r)$, and with $y = \sqrt{2}r$ we obtain $\dot y \le \sqrt{2}\Phi_a - \lambda_a \sin y$ on $y \le \pi$. This is Inequality (6.2) with $W + 2E$ replaced by $\sqrt{2}\Phi_a$ and κn replaced by λ_a, and the invariance and attraction conclusions follow as there. □

Proposition 13.1 (Weighted complete graphs, sup norm)**.** *Suppose the internal weights satisfy $K^a_{ij} \ge \kappa_{\min} > 0$ for all $i \ne j$, not necessarily symmetric. Then Theorem 6.1 holds with κ replaced by $\kappa_{\min}$ and E the largest external coupling of any member.*

Proof. In the proof of Theorem 6.1, the coupling terms for the extreme members are

$$-\sum_{j\ne p} K_{pj}\sin a_j - \sum_{j\ne q} K_{qj}\sin b_j.$$

All sines are nonnegative for $D \le \pi$, so the terms are at most

$$-\kappa_{\min}\sum_j(\sin a_j + \sin b_j) \le -\kappa_{\min} n \sin D,$$

using $a_p = b_q = 0$ for the omitted terms. □

Remark 13.1 (The uniform case)*.* For uniform all-to-all coupling κ the Laplacian has $\lambda_a = \kappa n$, and Proposition 13.1 returns Theorem 6.1 exactly. Theorem 13.1 returns the same coupling margin κn but measures dispersion in the Euclidean norm Ω_a and external push in $\|\bar d^{\,a}\|_2$ instead of in the sup norm, and converts back to the spread with a factor $\sqrt{2}$. For dense graphs the sup-norm version is sharper; for sparse graphs only the graph version applies, and its margin is the algebraic connectivity, so a weak bridge inside a group enters the condition through a small λ_a.

Remark 13.2 (Relation to known synchronization conditions)*.* Bounds of the form dispersion over coupling margin, and conditions stated through the algebraic connectivity of the interconnection graph, are standard in the synchronization literature on networks of phase oscillators [2], which treats phase cohesiveness and graph-dependent conditions for homogeneous and heterogeneous oscillators on complete and sparse graphs. The cohesion theorems of this chapter are not offered as new methods in that sense, and no comparison is made here with the sharpest known conditions in that survey. What is specific to this monograph is what is built on them: the bounded effect of a changing external coupling, release and recapture with explicit error accounting through the reduction theorem, the box of admissible adaptation, and the propagation of remainders through levels.

13.3 Reduction and the Principal Theorem for Nonuniform Coupling

Define the collective phase $\bar\vartheta_a$ as the mean of the lifted phases, as before, and let $S(t) = \mu(t)\bar S$.

Theorem 13.2 (Nonuniform reduction)**.** *With symmetric internal weights,*

$$\dot{\bar{\vartheta}}_g = \bar{f}_g + (S/n_g)\sin u + \rho_g,$$
$$\dot{\bar{\vartheta}}_h = \bar{f}_h - (S/n_h)\sin u + \rho_h,$$

with $u = \bar{\vartheta}_h - \bar{\vartheta}_g - a_{gh}$, $|\rho_g| \le (S/n_g)(D_g + D_h)$ *and* $|\rho_h| \le (S/n_h)(D_g + D_h)$. *Consequently* $\dot{u} = \Delta\bar{\omega} - N(t)\sin u + r(t)$, *where* $N = S(1/n_g + 1/n_h)$ *and* $|r| \le N(D_g + D_h)$.

Proof. Average the oscillator equations over each group. The internal terms cancel by symmetry. For the cross terms replace ϑ_j by $\bar{\vartheta}_h$ and ϑ_i by $\bar{\vartheta}_g$; each such replacement changes the argument of the sine by at most $D_g + D_h$, and the weights sum to S over the pairs. Subtracting the two equations gives the relative phase equation. □

Theorem 13.3 (Hierarchical release and recapture, nonuniform two-group case)**.** *Under the setting of this chapter assume* $\sqrt{2}\Phi_a < \lambda_a$ *for* $a = g, h$, *define* γ_a *as in Theorem 13.1, and assume* $s^* = |\Delta\bar{\omega}|/\bar{N} + \gamma_g + \gamma_h < 1$, *with* $\gamma_* = \arcsin s^*$. *Suppose* $\|e_a(0)\|_2 \le \gamma_a/\sqrt{2}$. *Then for every control* $\mu(t) \in [0,1]$ *the spreads satisfy* $D_a(t) \le \gamma_a$. *If* μ *is held at a value* μ_{rel} *with* $\mu_{\mathrm{rel}}\bar{N}(1 + \gamma_g + \gamma_h) < |\Delta\bar{\omega}|$ *the relative phase drifts monotonically at a rate bounded below by* $|\Delta\bar{\omega}| - \mu_{\mathrm{rel}}\bar{N}(1 + \gamma_g + \gamma_h)$. *If* μ *is restored to* 1 *at a time* t_r *at which* $|u(t_r)| < \pi - \gamma_*$ *in the relative-phase sense, then* $\limsup |u| \le \gamma_*$, *and for every* $\eta > 0$ *with* $\gamma_* + \eta < |u(t_r)|$ *the relative phase enters and stays in* $[-\gamma_* - \eta, \gamma_* + \eta]$ *after a finite time as in Theorem 9.1, with* $c_\eta = \bar{N}(\sin(\gamma_* + \eta) - s^*)$ *giving the uniform bound* π/c_η *for* $|u(t_r)| \le \pi - \gamma_* - \eta$. *Every cross pair then satisfies* $\limsup |\vartheta_j - \vartheta_i - a_{gh}| \le \gamma_* + \gamma_g + \gamma_h$ *modulo* 2π.

Proof. The external perturbation on member i is at most $\mu(t)\bar{d}_i \le \bar{d}_i$, so $\|\varepsilon\|_2 \le \|\bar{d}^a\|_2$ and Theorem 13.1 gives the spread bounds for every control. Theorem 13.2 with $D_a \le \gamma_a$ gives $|r| \le N(t)(\gamma_g + \gamma_h)$. The release and recapture statements then follow exactly as in the proof of Theorem 12.1, with $\bar{N}$ in place of the uniform value. □

Theorem 12.1 is the special case of uniform weights on complete graphs, since then $\lambda_a = \kappa_a n_a$, $\bar{S} = n_g n_h \bar{\kappa}$ and $\bar{N} = (n_g + n_h)\bar{\kappa}$, up to the replacement of sup-norm by Euclidean dispersion noted above.

13.4 Amplitude and Offset Commands

The fish-robot controller of Wang and colleagues separates three layers: input saturation functions, coupled oscillators, and an output transition function to servo commands [22]. In their oscillator layer the amplitude a_i and offset x_i of oscillator i obey critically damped second-order dynamics toward commanded values A_i and X_i,

$$\ddot{a}_i = \alpha\big(\alpha(A_i - a_i) - 2\dot{a}_i\big), \qquad \ddot{x}_i = \beta\big(\beta(X_i - x_i) - 2\dot{x}_i\big), \tag{13.2}$$

with output $\theta_i = x_i + a_i\cos\phi_i$ and $\alpha = \beta = 20$ rad/s in their experiments. Their phase dynamics are second order, with coupling terms that carry the amplitude a_j of the sending oscillator, so in their model a change of commanded amplitude also changes the strength of the phase coupling. Equations (13.2) and the output form are theirs. What follows adapts the structure to tendon take-up and is an extension proposed here, not their mechanical implementation.

Adaptation to cords. Let A_i be the commanded cord excursion, B_i the commanded mean cord length and β_i the desired timing offset, and let

$$\ddot{a}_i + 2\lambda_a \dot{a}_i + \lambda_a^2 (a_i - A_i) = 0, \qquad \ddot{b}_i + 2\lambda_b \dot{b}_i + \lambda_b^2 (b_i - B_i) = 0,$$
$$\ell_i^{\mathrm{des}}(t) = b_i(t) + a_i(t) \cos\bigl(\theta_i(t) - \beta_i\bigr). \tag{13.3}$$

This is the structure of (13.2) with $\lambda_a = \alpha$ and $\lambda_b = \beta$. In the phase model of this monograph the amplitude and offset states do not enter the phase equations, which is a modeling choice different from theirs.

Proposition 13.2 (Amplitude transient)**.** *For constant commands the amplitude error is $a_i(t) - A_i = (c_1 + c_2 t) e^{-\lambda_a t}$ with $c_1 = a_i(0) - A_i$ and $c_2 = \dot{a}_i(0) + \lambda_a c_1$, and likewise for $b_i - B_i$. If $\dot{a}_i(0) = 0$ then $|a_i(t) - A_i| = |c_1|(1 + \lambda_a t) e^{-\lambda_a t}$ is nonincreasing, so the amplitude approaches its command without overshoot.*

Proof. $\lambda = \lambda_a$ is a double root of the characteristic polynomial $s^2 + 2\lambda s + \lambda^2$, which gives the form, and the constants follow from the initial values. For $\dot{a}(0) = 0$, $\frac{\mathrm{d}}{\mathrm{d}t}\bigl[(1 + \lambda t) e^{-\lambda t}\bigr] = -\lambda^2 t e^{-\lambda t} \le 0$. □

A change of commanded excursion or mean length therefore produces a smooth transient at the output while the phase-coupling network, and with it the cohesion and capture results of this part, is untouched, since those results bound phase spreads only. This is a second way of changing a movement, alongside release by decoupling and detuning, and it changes amplitude and offset and not timing. The drum then follows ℓ_i^{des} through the first-order lag of Chapter 31, so the command is shaped by the critically damped stage before the drum sees it.

Chapter 14

Adaptation of Internal Couplings

Practice and fatigue are changes to the couplings themselves. This chapter proves that, in the nonuniform two-group setting of Chapter 13, the cohesion, release and recapture results survive adaptation of selected internal weights, provided the adaptation keeps the weights inside an admissible region defined from the theorem inequalities. Slowness of adaptation is not what protects the result. A slow law that can drive a weight below the certified floor eventually crosses the threshold, and a fast law that cannot do so never does.

14.1 Admissible Region

Keep the setting of Chapter 13: fixed internal topology, fixed target offsets, fixed cross weights $\bar{K}_{ij}$ with the control $\mu(t)$. Split the internal edges of group a into fixed edges, whose weights do not change, and adaptive edges e, whose weights $k_e(t)$ are confined to an interval $[\ell_e, h_e]$ with $0 \le \ell_e < h_e$. Let L_a^{lo} be the Laplacian of group a with every adaptive edge at its floor ℓ_e, and let $\underline{\lambda}_a$ be its algebraic connectivity.

Lemma 14.1 (Floor certificate)**.** *If $\ell_e \le k_e \le h_e$ for every adaptive edge, the Laplacian $L_a(k)$ of group a satisfies $\lambda_2(L_a(k)) \ge \underline{\lambda}_a$.*

Proof. $L_a(k) - L_a^{\mathrm{lo}} = \sum_e (k_e - \ell_e) L_e$ with each edge Laplacian $L_e \succeq 0$, so $x^\top L_a(k) x \ge x^\top L_a^{\mathrm{lo}} x$ for all x. The Courant-Fischer characterization $\lambda_2 = \min\{x^\top L x : \|x\| = 1,\ x \perp \mathbf{1}\}$ then gives the claim. □

The admissible region is the set of weight vectors in the box $\prod_e [\ell_e, h_e]$, and it is certified when the floor data satisfy the hypotheses of Theorem 13.3 with $\underline{\lambda}_a$ in place of λ_a: that is, $\sqrt{2}\Phi_a < \underline{\lambda}_a$, $\underline{\gamma}_a = \arcsin(\sqrt{2}\Phi_a / \underline{\lambda}_a)$, and $\underline{s} = |\Delta\bar{\omega}|/\bar{N} + \underline{\gamma}_g + \underline{\gamma}_h < 1$. Because Φ_a does not involve the internal weights, only $\underline{\lambda}_a$ enters.

14.2 A Law That Preserves the Region

A law such as $\dot{K}_{ij} = \epsilon(\eta \cos(\vartheta_j - \vartheta_i) - \lambda K_{ij})$ does not preserve nonnegativity: at $K_{ij} = 0$ the derivative is negative whenever the cosine is, so the constraint must be enforced by the law

and not merely written beside it. Parametrize each adaptive weight as $k_e = \ell_e + (h_e - \ell_e)\sigma_e$ and adapt the gate variable by

$$\dot{\sigma}_e = \sigma_e(1-\sigma_e)\, G_e(t,x), \qquad \sigma_e(0) \in [0,1], \tag{14.1}$$

where x collects the phases and gate variables and G_e is any function, continuous in t and locally Lipschitz in x, that may depend on the state arbitrarily and carries the rate of adaptation. A practice and fatigue law is $G_e = \epsilon_e\big[\eta\cos(\vartheta_j - \vartheta_i) - \phi_e(t)\big]$ with a fatigue load $\phi_e \geq 0$: pairs near their preferred offset strengthen, pairs far from it weaken, and accumulated load weakens every edge.

Lemma 14.2 (Invariance of the box)**.** *For every admissible G_e, $\sigma_e(t) \in [0,1]$ for all $t \geq 0$, hence $\ell_e \leq k_e(t) \leq h_e$.*

Proof. The vector field is locally Lipschitz; on any interval of existence σ_e stays bounded once the argument below is applied, and the phase derivatives are bounded because the weights lie in the box, so the state cannot blow up in finite time and solutions exist for all time. Along a solution $g(t) = G_e(t, x(t))$ is a continuous function of time, and $\dot{\sigma} = \sigma(1-\sigma)g(t)$ has the constant solutions $\sigma \equiv 0$ and $\sigma \equiv 1$. By uniqueness a solution that starts in $[0,1]$ cannot reach either constant in finite time without coinciding with it, and therefore stays in $[0,1]$. □

14.3 The Adaptation Theorem

Theorem 14.1 (Coordination under adaptation)**.** *Assume the floor data are certified as above, adaptive weights obey* (14.1) *with arbitrary G_e, internal weights remain symmetric, and $\|e_a(0)\|_2 \leq \underline{\gamma}_a/\sqrt{2}$. Then for every adaptation rate, every control $\mu(t) \in [0,1]$ and every state dependence of G_e, the spreads satisfy $D_a(t) \leq \underline{\gamma}_a$ for all t, and the release and recapture statements of Theorem 13.3 hold with $\underline{\gamma}_a$ in place of γ_a and $\underline{s}$ in place of s^*.*

Proof. By Lemma 14.2 and Lemma 14.1, $\lambda_2(L_a(t)) \geq \underline{\lambda}_a$ at every time. The proof of Theorem 13.1 uses the Laplacian only through the pointwise inequality

$$e^\top L_c(t) e \geq \operatorname{sinc}(\sqrt{2}r)\lambda_2(L_a(t))r^2$$

and the symmetry of the weights, so it holds for time-varying weights, with λ_a replaced by the uniform lower bound $\underline{\lambda}_a$. Theorem 13.2 is an exact identity for symmetric time-varying internal weights, since the internal terms cancel at every instant, and the cross weights are unchanged. The release and recapture arguments depend only on the reduced equation and the spread bounds, which are now those with $\underline{\gamma}_a$. □

The theorem is a statement about a forward-invariant region, and it separates two questions that are easy to run together. Slowness is not used anywhere in the proof, and the result holds for adaptation as fast as one likes; what protects the coordination is that the weights cannot leave the certified box. Conversely, slowness cannot replace the floor.

Proposition 14.1 (Slowness without a floor)**.** *Consider two strongly coupled clusters of sizes n_1, n_2 joined by a single adaptive edge of weight $b(t)$ with floor $\ell = 0$, and let the*

clusters have mean frequency difference $\Delta\omega > 0$. In the reduced description the cluster phase difference obeys $\dot{\Phi} = \Delta\omega - b(t)(1/n_1 + 1/n_2)\sin\Phi$, which has no locked solution whenever $b(t) < \Delta\omega/(1/n_1 + 1/n_2)$. If the gate dynamics (14.1) drive $\sigma \to 0$ with $G \le -g_0 < 0$, then $b(t) \to \ell = 0$ and this condition holds after a finite time, however small ϵ is.

This is an argument in the reduced model, not a proof for the full system, and it is included to show that the floor requirement cannot be weakened to a rate restriction. Simulation shows the same behavior in the full system (Chapter 16).

14.4 Moving Targets and Drifting Frequencies

Let the target offsets move, $\beta_i(t)$ and $a_{gh}(t)$ of class C^1, with the preferred offsets following them, $\alpha_{ij}(t) = \beta_j(t) - \beta_i(t)$ within a group and $\alpha_{ij}(t) = \beta_j(t) - \beta_i(t) + a_{gh}(t)$ across groups. With $\vartheta_i = \theta_i - \beta_i(t)$ the phase equation is unchanged except that the effective frequency becomes $f_i - \dot{\beta}_i$, and the relative phase $u = \bar{\vartheta}_h - \bar{\vartheta}_g - a_{gh}(t)$ obeys the same equation with $\Delta\bar{\omega}$ replaced by $\Delta\bar{\omega} - \dot{a}_{gh}(t) + \overline{\dot{\beta}}_g - \overline{\dot{\beta}}_h$ where bars denote group means.

Corollary 14.1 (Moving targets)**.** *Theorems 13.1, 13.2, 13.3 and 14.1 hold with Ω_a computed from $\sup_t \|(f - \dot{\beta}) - \text{mean}\|_{2,a}$ and with $\Delta\bar{\omega}$ replaced by the effective mismatch above, provided the derivatives $\dot{\beta}_i$ and $\dot{a}_{gh}$ are bounded, so that the suprema are finite. The effective mismatch may now depend on time. The capture and spread arguments use only a pointwise upper bound on $|\Delta\bar{\omega}_{\text{eff}}(t)|$, which is then the quantity entering s^*; the release statement needs, over the interval considered, a positive lower bound on $|\Delta\bar{\omega}_{\text{eff}}(t)|$ together with a constant sign, since monotone drift of u depends on the sign. The initial-condition hypothesis on $\|e_a(0)\|_2$ is unchanged, and the phases are those of the shifted variables $\vartheta_i = \theta_i - \beta_i(t)$.*

The rate of change of the target therefore enters the disturbance estimate explicitly, through the dispersion of $\dot{\beta}_i$ within a group and through $\dot{a}_{gh}$ across groups. Time-dependent natural frequencies, as in fatigue-induced drift of ω_i, are already covered, since Ω_a is defined as a supremum over time of the deviation of $f_i(t)$.

Chapter 15

Three-Level Hierarchies

15.1 Setting

Take two super-groups A and B, each consisting of m groups of n oscillators, with uniform couplings κ_1 between oscillators of the same group, κ_2 between oscillators of different groups of the same super-group, and $\kappa_3(t) = \mu(t)\bar{\kappa}_3$ between oscillators of different super-groups, $\mu \in [0,1]$. Offsets are consistent within each super-group, so that the shift lemma removes them, and the offset between super-groups is a_{AB}. All frequency spreads are suprema over time, and the mismatch $\Delta\bar{\omega}$ is constant. Let W_1 be the largest spread of effective oscillator frequencies within any group, W_2 the largest spread of group-mean effective frequencies within a super-group, and $\Delta\bar{\omega}$ the difference of the super-group mean frequencies.

15.2 Propagation of the Remainders

Define in turn

$$E_1 = (m-1)n\kappa_2 + mn\bar{\kappa}_3, \qquad \gamma_1 = \arcsin\frac{W_1 + 2E_1}{\kappa_1 n}, \tag{15.1}$$

$$E_2 = mn\bar{\kappa}_3 + 2\gamma_1 E_1, \qquad \gamma_2 = \arcsin\frac{W_2 + 2E_2}{mn\kappa_2}, \tag{15.2}$$

$$D_3 = 2\gamma_1 + \gamma_2, \qquad \bar{N}_3 = 2mn\bar{\kappa}_3, \qquad s^* = \frac{|\Delta\bar{\omega}|}{\bar{N}_3} + 2D_3. \tag{15.3}$$

Theorem 15.1 (Three-level release and recapture)**.** *Assume the arguments of both arcsines are less than one, that $s^* < 1$, and that every group has spread at most γ_1 and every super-group has group-collective-phase spread at most γ_2 at time zero. Then for every control $\mu(t)$: (a) every group has spread at most γ_1 for all time; (b) the collective phases of the groups within each super-group have spread at most γ_2 for all time; (c) all oscillators of each super-group lie within a spread D_3 for all time; (d) if μ is held at μ_{rel} with $2mn\,\mu_{\mathrm{rel}}\bar{\kappa}_3(1+2D_3) < |\Delta\bar{\omega}|$, the relative phase of the super-groups drifts monotonically; (e) if μ is restored to one when the relative phase satisfies $|U| < \pi - \gamma_*$ with $\gamma_* = \arcsin s^*$, then $\limsup|U| \le \gamma_*$, and for every $\eta > 0$ with $\gamma_* + \eta < |U|$ at the restoration time U enters and stays in $[-\gamma_* - \eta, \gamma_* + \eta]$ after a finite time as in Theorem 9.1.*

Proof. (a) An oscillator in a group has external coupling at most $(m-1)n\kappa_2+mn\bar{\kappa}_3 = E_1$, so Theorem 6.1 applies with $\kappa = \kappa_1$ and $W = W_1$. (b) By Theorem 7.1 the collective phases $\bar{\vartheta}_a$ of the groups of A obey an equation with uniform symmetric coupling $n\kappa_2$ among m effective oscillators, effective frequencies $\bar{\omega}_a$ with spread at most W_2, and a disturbance made of two parts: the coupling to the groups of B, at most $mn\bar{\kappa}_3$ in modulus, and the remainder ρ_a, at most $\sum_b n\kappa_{ab}(D_a + D_b) \le 2\gamma_1 E_1$ by (7.2) and part (a). So the disturbance is bounded by E_2, and Theorem 6.1 applies with coupling $n\kappa_2$ and group size m. (c) For i in group a and j in group b of the same super-group, $|\vartheta_i - \vartheta_j| \le |\vartheta_i - \bar{\vartheta}_a| + |\bar{\vartheta}_a - \bar{\vartheta}_b| + |\bar{\vartheta}_b - \vartheta_j| \le \gamma_1 + \gamma_2 + \gamma_1$, because a member's deviation from its group mean is at most the group spread. (d) and (e) Let $\bar{\vartheta}_A, \bar{\vartheta}_B$ be the means of all lifted phases in the super-groups, which equal the means of the group means, and let $U = \bar{\vartheta}_B - \bar{\vartheta}_A - a_{AB}$. Applying Theorem 13.2 at the oscillator level to A and B, whose internal weights are symmetric and whose spreads are at most D_3, gives $\dot{U} = \Delta\bar{\omega} - N_3(t)\sin U + r$ with $N_3 = 2mn\mu\bar{\kappa}_3$ and $|r| \le N_3 \cdot 2D_3$. The release and capture arguments of Theorem 12.1 then apply with $\gamma_g + \gamma_h$ replaced by $2D_3$. □

15.3 Does the Error Estimate Stay Useful?

The estimate grows with the number of levels. The remainder from level one enters the level-two disturbance as $2\gamma_1 E_1$, so

$$\sin\gamma_2 = \frac{W_2 + 2mn\bar{\kappa}_3}{mn\kappa_2} + 4\gamma_1\Big(\frac{m-1}{m} + \frac{\bar{\kappa}_3}{\kappa_2}\Big), \tag{15.4}$$

and the second term is the contribution of the lower level, roughly $4\gamma_1(m-1)/m$, which is between $2.7\gamma_1$ for $m = 3$ and $4\gamma_1$ for large m. The final error budget is then $2D_3 = 4\gamma_1 + 2\gamma_2 \approx (4 + 2\cdot 4(m-1)/m)\gamma_1 + \ldots$, which is about nine times the level-one margin for $m = 3$ before the frequency terms are added. A third level therefore demands that γ_1 be of order 0.05 or smaller for the condition $s^* < 1$ to hold, which means strong separation of scales: in the example below $\kappa_1/\kappa_2 = 100$ and $\kappa_2/\bar{\kappa}_3 = 20$. In this estimate each additional level multiplies the requirement by a similar factor, so it is useful only for shallow hierarchies with strongly separated couplings and degrades geometrically with depth. This is a property of the present estimate, which treats each remainder as a worst-case disturbance at the next level. It has not been established as an intrinsic limitation of hierarchical coordination, and on the same example the direct bound of the next remark and the observed spreads of Chapter 16 are smaller.

Remark 15.1 (A direct route). Theorem 13.1 applied to a super-group as a single group, with its full internal weight graph, avoids the propagation. For the parameters of the example below the direct bound on the full spread of a super-group is 0.236, against the propagated $D_3 = 0.319$. The direct route needs the algebraic connectivity of the whole super-group graph, which is $mn\kappa_2$ for strongly separated block structure, but it does not expose the hierarchy, and the propagated version remains the one that shows how each level's error enters the next.

Chapter 16

Boundary Simulations

The simulations below integrate the oscillator equations by a fourth-order Runge-Kutta scheme and test the proved bounds and their boundaries. They are consistency checks and do not substitute for proofs; scripts accompany the draft.

16.1 Nonuniform Two-Group Case

Groups of six and five oscillators with sparse internal rings plus one chord, internal weights between 3 and 6, and sparse random cross weights up to 0.3 were used. The theorem conditions held with $\gamma_g = 0.338$, $\gamma_h = 0.299$, $s^* = 0.679$ and $\gamma_* = 0.747$. Over a run with locking, a release of fifty time units, and recapture, the largest observed spreads were 0.159 and 0.139, a factor of two to three below the bounds. After recapture the relative phase settled at -0.045, inside $[-\gamma_*, \gamma_*]$.

16.2 Initial Phases Near the Excluded Band

Recapture was started with the relative phase set to 0.5, 2.0, 2.345, 2.918, π and -3.04. The last three lie inside the excluded band around the antipode of half-width γ_*, and 2.345 lies just outside it. All six runs ended at the same relative phase of -0.045. These runs show that the sufficient condition is conservative at this parameter set. They do not show that the excluded band lacks dynamical significance. In the remainder-free equation the band contains the unstable equilibrium $\pi - a$; here $a = -0.043$, so it sits at $\pi + 0.043$, close to but not equal to the tested value π, and a trajectory started exactly on it would not be captured. An admissible remainder can move that equilibrium anywhere within the band, which is why the sufficient condition excludes the band. Six initial phases at one parameter set cannot rule out failures at other admissible remainders or parameters.

16.3 Asymmetric Forcing and Weak Internal Connections

A constant drive was added to one oscillator of the first group. The proved condition fails for drives between 2 and 4, while the observed spread grew slowly.

Drive on one member	$\sqrt{2}\Phi_g/\lambda_g$	Proved bound γ_g	Observed maximum spread
0	0.331	0.338	0.119
0.5	0.418	0.431	0.119
1	0.547	0.579	0.119
2	0.815	0.953	0.174
4	1.356	none	0.360
8	2.439	none	0.748

A weak internal connection was tested by splitting a group of six into two triangles with internal weight 5 joined by a single bridge of weight b, with the two triangles at frequencies differing by 0.15, and with the external coupling switched off so that the test measured the group alone.

Bridge weight b	λ_2	Proved bound	Observed maximum spread
4	1.895	0.141	0.119
2	1.118	0.240	0.148
1	0.610	0.451	0.262
0.5	0.319	0.986	0.502
0.3	0.195	none	0.883
0.25	0.163	none	1.153
0.2	0.131	none	3.516
0.1	0.066	none	9.874

The proved condition ceases to hold at a bridge weight near 0.4. The group actually lost cohesion between $b = 0.25$ and $b = 0.2$, which is close to the two-cluster locking threshold $b \geq 3\Delta\omega/2 = 0.225$. The proved range of applicability is therefore conservative by a factor of about two in the bridge weight, and the bound itself is conservative by about a factor of two in the spread where it applies. With the external coupling switched on, the same group stayed cohesive down to $b = 0.1$, because the cohesive neighboring group acted as a common reference. This is the effect described in Remark 6.1.

16.4 Adaptation

The two-group setting of the first section was run with all fourteen internal weights adaptive, floors at 0.8 times and ceilings at 1.6 times the original weights, which certifies $\underline{\lambda}_g = 3.81$ and $\underline{\lambda}_h = 4.83$, $\underline{\gamma}_g = 0.427$, $\underline{\gamma}_h = 0.377$ and $\underline{\gamma}_* = 1.010$. Release lasted from time 40 to 90, after which the relative phase was set to 1.5.

Adaptation law	Minimum λ_2	Maximum spread	Final u
Practice with mild fatigue, rate 0.05	4.76, 6.04	0.201, 0.175	−0.044
Practice, rate 5	4.77, 6.05	0.201, 0.175	−0.044
Adversarial push to floor, rate 5	3.81, 4.83	0.201, 0.175	−0.045
Adversarial push to ceiling, rate 5	4.78, 6.06	0.201, 0.175	−0.044
Alternating sign (period 7), rate 5	4.78, 6.06	0.201, 0.175	−0.044
No floor, push to zero, rate 1	0.00, 0.00	13.35, 13.44	1.57

In the first five runs the weights obeyed the box, the minimum λ_2 never went below the certified floor, and the spreads stayed inside the bounds, at every rate tested including the adversarial pushes. The maximum spread in those runs equals the initial spread, so they test invariance and not the tightness of the bound. In the last run the floor was removed and the law pushed every weight to zero; cohesion was lost, which the theorem permits once the floor certificate is gone and which the proposition on slowness without a floor illustrates in the reduced model; no recapture guarantee applies.

16.5 Three Levels

With $n = m = 3$, $\kappa_1 = 100$, $\kappa_2 = 1$, $\bar{\kappa}_3 = 0.05$, $W_1 = 0.1$, $W_2 = 0.05$ and $\Delta\bar{\omega} = 0.05$, the propagated bounds were $\gamma_1 = 0.043$, $\gamma_2 = 0.232$, $D_3 = 0.319$ and $s^* = 0.693$. The largest observed spreads were 0.009, 0.111 and 0.120 respectively, smaller by factors of about 4.6, 2.1 and 2.7. Release over fifty time units made the relative phase of the super-groups drift, and after recapture it returned to 0.056, inside $[-\gamma_*, \gamma_*]$ with $\gamma_* = 0.765$. The bounds held in every run. The finding about depth is a statement about the present estimate, which requires strongly separated couplings, and the observed spreads were well inside it.

Part IV

Summaries Between Levels

Chapter 17

Summaries Passed Between Levels

Scope of this part. Established here are the summary Z_a, the two-sided relation between coherence and spread, the distinction between composition and aggregation, and a conditional proposition on resets. The identification of Z_a with a motion-generating amplitwist, and the compatibility of amplitwist composition with hierarchical reduction, are open interpretations and nothing in Parts II and III depends on them.

Each level of the hierarchy reports to the next a summary of its coherence and phase. In the notation of Part II this summary is the pattern-adjusted order parameter $Z_a = r_a e^{\mathrm{i}\Psi_a}$ of Equation (4.3), whose modulus is an amplitude and whose argument is a rotation. The author's essay on locomotion [25] calls such a pair an amplitwist, a reading this monograph treats as an open interpretation (Section 17.4), borrowing the polar decomposition of the derivative of a holomorphic map, $f'(z_0) = re^{\mathrm{i}\theta}$ with $r = |f'(z_0)|$ and $\theta = \arg f'(z_0)$. It defines an amplitwist operator along a curve as a local, instantaneous description $A(t) = r(t)e^{\mathrm{i}\theta(t)}$ of the first-order behavior of the generating process, and an amplitwist chain as the composition $A^{(k+1)} = C(A_1^{(k)}, \dots, A_{n_k}^{(k)})$ of a coupling functional C over levels from the finest actuator to the coarsest whole-body gait.

17.1 What the Results of Part III Supply

Two ingredients of that proposal are made precise by the theorems of this monograph. First, the claim that coherence tends to one under strong coupling becomes a quantitative statement: by Theorem 6.1 and Theorem 13.1 a group whose spread satisfies the cohesion bound has $r_a \geq \cos\gamma_a$ by Lemma 4.2, with γ_a given explicitly by the coupling margin and the dispersion. Second, the coupling functional C has an explicit realization in the Kuramoto setting. Theorem 7.1 shows that the collective phases of the groups obey

$$\dot{\bar{\vartheta}}_a = \bar{\omega}_a + \sum_{b \neq a} n_b \kappa_{ab} \sin(\bar{\vartheta}_b - \bar{\vartheta}_a - a_{ab}) + \rho_a, \qquad |\rho_a| \leq \sum_{b \neq a} n_b \kappa_{ab} (D_a + D_b),$$

which is a coupling functional acting on the level-k summaries $(\bar{\vartheta}_b, D_b)$ and delivering the level-$(k+1)$ phase velocity, with a bounded remainder. The three-level Theorem 15.1 is the statement for a chain with two applications of this functional, and its error propagation is the quantitative behavior of the chain under the present estimate. The summary of Part II uses

the argument Ψ_a of the order parameter, whereas Theorem 7.1 is an identity for the mean phase $\bar{\vartheta}_a$; the two are different quantities, and they differ by at most the spread D_a when the group lies in an arc narrower than π. The coherence bound above is unaffected, since Lemma 4.2 holds for any reference angle. Reading the reduction equation as an equation for the phase of the amplitwist Z_a would require carrying that additional error D_a, and no such statement is proved here.

17.2 What the Summary Cannot See

Distinct internal configurations can have identical summaries. Take a group of four with target offsets removed and phase errors

$$(+\epsilon, -\epsilon, +\epsilon, -\epsilon) \quad \text{and} \quad (+\epsilon, +\epsilon, -\epsilon, -\epsilon)$$

in the first and second configurations. Both have $Z = \cos\epsilon$, the same coherence and phase, yet they differ in which members lead. If the connection to the higher level weights the members equally the difference is invisible upward; if it weights some members more heavily, as when only two members are coupled to the next level, the two configurations send different signals. The sufficiency of the summary is therefore a property of the upward connection and not of the group alone, and any theorem of hierarchical reduction must state the class of connections for which the summary suffices. Theorem 7.1 does so for uniform cross coupling, where only the group mean and the spread enter.

17.3 Composition and Aggregation

Composition of operators and aggregation of populations are different operations.

Proposition 17.1. *Composition satisfies $A_a A_b(z) = Z_a Z_b z$, with modulus $r_a r_b$ and argument $\Psi_a + \Psi_b$. Aggregation of two groups of sizes n_a, n_b gives $Z = (n_a Z_a + n_b Z_b)/(n_a + n_b)$. These coincide only in special cases.*

Proof. For $r_a = r_b = 1$, $\Psi_a = 0$, $\Psi_b = \pi/2$, $n_a = n_b$, the composition has modulus one and the aggregate has modulus $\sqrt{2}/2$. □

A compatibility result between amplitwist composition and hierarchical reduction is needed wherever the monograph identifies the two, and it is listed as open. The theorems of Parts II and III do not depend on it.

17.4 The Open Interpretation

What remains unproved is the geometric reading. Writing $Z_a = r_a e^{\mathrm{i}\Psi_a}$ establishes a polar representation of the summary. The further claim that Z_a is the derivative of a map generating bodily motion needs a defined space on which the map acts and an argument that its derivative at the point of interest is Z_a. For the linear map $f(z) = Z_a z$ the derivative is trivially Z_a, so the claim has content only for a specified, nonlinear generating process, and no such process is given here. The ledger records this as open. The established content of this

chapter does not depend on it: the summary, the reduced equation, the two-sided relation between coherence and spread, and the distinction between composition and aggregation hold whether or not the summary has a geometric reading. The amplitwist is therefore presented as an interpretation proposed in the author's essay and left open here.

Chapter 18

Chains of Joints, Limbs and Gaits

18.1 Levels and Assembly Index

The locomotion essay organizes a body into joint, limb, gait and whole-body levels, each obtained from the previous by one application of the coupling functional, with the gait operator coupled to postural and vestibular loops. It defines the assembly index of a chain as $n - 1$, the number of coupling applications, and notes that it is bounded below by the number of anatomical levels and does not depend on how many oscillators lie at each level [25]. The theorems of this monograph fit that organization: Theorem 13.3 is a chain with assembly index one, and Theorem 15.1 one with index two. The bounds do depend on the number of oscillators through the coupling margins κn and the algebraic connectivity, so the independence of the assembly index from oscillator count is a statement about depth and not about the size of the error estimates.

18.2 The Bipedal Case

For two legs the essay writes the gait operator as $A^{(2)} = C_K(A_L^{(1)}, A_R^{(1)})$ and notes that an anti-phase gait, with preferred offset $\Delta\phi_{LR} \approx \pi$, would appear as low coherence under the unadjusted order parameter, so that it uses the offset-corrected form $\tilde{Z} = \frac{1}{2}\big(e^{\mathrm{i}\theta_L} + e^{\mathrm{i}(\theta_R - \pi)}\big)$. This is the pattern-adjusted order parameter of Equation (4.3) with target offsets $\beta_L = 0$, $\beta_R = \pi$, and the Shift Lemma shows that the offset-corrected gait dynamics are the in-phase dynamics in shifted variables. The generalization made here is to arbitrary offsets β_i for every oscillator in every group, with the cohesion bound stated for the shifted spread.

18.3 Contact Events

The essay observes that the Kuramoto coupling is smooth, so the summary Z is well defined between contact events, and it proposes, as a conjecture limited to planar single-point contact, that the reset at foot strike is a single amplitwist operator, $Z(t^{*+}) = A_e Z(t^{*-})$ with $A_e = \rho_e e^{\mathrm{i}\phi_e}$. It describes the chain as currently a theory of swing phase. A conditional result for general resets follows from the cohesion theorems.

Proposition 18.1 (Invariance across resets)**.** *Consider the flow of Theorem 13.1 or Theorem 6.1, interrupted at event times $t_1 < t_2 < \dots$ by continuous reset maps R_k on the phases. Assume (i) the hypotheses of the cohesion theorem hold between events, (ii) each R_k maps the set $\{D_a \le \gamma_a \text{ for all } a\}$ into itself, and (iii) the events do not accumulate in finite time, for instance because consecutive events are separated by a positive dwell time. Then the set is forward invariant for the hybrid system.*

Proof. Between events the set is invariant by the cohesion theorem. At an event the state lies in the set and is mapped into it by (ii). Condition (iii) means the time axis is exhausted by finitely many intervals on every bounded time interval, so induction over the events covers all time. □

For the essay's conjecture the proposition has a concrete reading, but a reduction of the magnitude of the summary does not by itself say how much the phase spread widens, and that relationship needs its own bound. Lemma 4.3 supplies one. Suppose a reset scales $|Z|$ by $\rho < 1$ and that, before the event, the state lies in the inner set $\{D_a \le \gamma_{\text{in}}\}$ with $\gamma_{\text{in}} < \gamma_a$, so that $r \ge \cos(\gamma_{\text{in}}/2)$. Then after the reset $r' \ge \rho\cos(\gamma_{\text{in}}/2)$, and the phases lie in an arc of width at most $2\arccos\big(1 - n(1 - r')\big)$ on the circle, which bounds the lifted spread D' provided the reset assigns lifts inside that arc (a reset that moves a phase by a full turn in the lift would not be covered). Condition (ii) of the proposition is therefore implied by

$$n\big(1 - \rho\cos(\gamma_{\text{in}}/2)\big) \le 1 - \cos(\gamma_a/2). \tag{18.1}$$

For example, with $n = 6$, $\gamma_a = 0.4$ and $\gamma_{\text{in}} = 0.1$ the condition requires $\rho \ge 0.998$. The bound is a worst-case guarantee, since it allows the reset to produce a single outlier phase, and it is very restrictive. A reset known to widen all phases symmetrically needs a separate and weaker condition, and the form of the reset has to be established for the particular contact. Condition (18.1) is a sufficient guarantee and not a necessary physical requirement: a body that violates it can still preserve cohesion if its actual resets have more structure than the worst case allows. It is a requirement on the body only in the sense of a sufficient design margin, and it does not prove the conjecture about the form of the reset. The idealized pneumatic chamber of Chapter 2 satisfies the non-accumulation condition by the dwell-time bound of its theorem, and its reset is the identity on pressure.

Part V

Endomarionettes

Chapter 19

The Endomarionette Body

Scope of this part. This part mixes categories and labels them where they occur. The body, spine, first embodiment and fabrication chapters are proposals. The derivations of the period, the interaction functions, the interface error budget and the cord-count bound are approximations or theorems in the stated models. The worked networks, the second cycle and the step-size study are numerical findings. The validation program and the bench protocol are specifications of tests that have not been run.

An endomarionette is a hollow paper or paper-mâché body in which internal cords, drawn from a central spinal apparatus, pull against flexible joints, with air pressure supplying or regulating actuation. The design has been described in the author's earlier writing under the names pneumatic endomarionettes and endo-marionettes, with a hexagonal cell structure, tubing and pressure as the pneumatic elements. A recent statement of the lightweight construction, which situates it among tendon-driven mechanisms, origami fabrication and soft robotics while leaving the arrangement of strings, bladders and valves unspecified, appears in [27].

19.1 Three Mechanical Roles

The mechanism has three parts that can look alike in a drawing and do different work. The spinal actuator takes up or pays out a length of cord. The cord-routing guides, called here internal guide knuckles, are fixed points inside the shell through which a cord passes; they set the path of the cord and the places where its pulling direction changes, and they transmit no net pull to the body except through the reaction of the bend. The attachment points are where the cord terminates on a distal part of the limb and tension acts on the shell. The guide knuckle is taken to be a smooth eye on a support fixed to the shell, so the cord slides through it and does not roll over a pulley. A guide is passed through and an attachment is terminated at, and the analysis below treats them as different objects.

19.2 The Blackbox

The winding assembly that holds the take-up mechanisms is called the blackbox. The name is also a modeling commitment: the theory treats the blackbox through its interface only. Its inputs are commands to its winding channels, and its outputs are, for each cord, the length

taken up and the tension delivered. What mechanism produces those outputs, whether a motor, a pressure-driven drum or a pneumatic oscillator, is left open, and the results of Parts II and III apply to any blackbox whose channel outputs satisfy the oscillator model.

19.3 Embodiments of the Blackbox

Because the blackbox is specified by its interface, several embodiments can realize it, and they differ in how the surplus cord is stored. In the spool embodiment the cord winds onto a spool, and the length taken up for a rotation φ is $s = \int_0^{\varphi} R(\phi)\,\mathrm{d}\phi$, where R is the effective winding radius. A constant radius gives $s = R\varphi$. When successive turns stack on one another R grows with the amount wound, so the length per rotation is not constant and the map from command to length must be calibrated. In the drop embodiment the cord is pulled through a gate and the surplus falls into a cavity below the mechanism. The length taken up is then the length pulled, the stored surplus applies no tension of its own, and paying out is passive.

The two embodiments satisfy the same interface and differ in their departures from it. The spool holds the cord under control at all times and keeps the take-up length a monotone function of the command, at the cost of the radius effect and friction in the winding. The drop gives a simple and nearly constant relation between pull and length, but it depends on the orientation of the body relative to gravity, and surplus cord can tangle or jam the gate. These are design considerations that have not been tested. They enter the theory as the delay, friction and calibration errors treated in the robustness obligation of Appendix B.

19.4 Spinal Arrangements

Two arrangements of the actuator are considered. In the concentrated arrangement a single blackbox in the thorax holds the take-up channels for all cords. In the segmental arrangement a series of smaller blackboxes lies along the spine, occupying the places of vertebrae, each routing or taking up the cords of one body region, so that the spine becomes an articulated succession of tendon-control units. The segmental arrangement matches the hierarchical structure of Part II, since each unit can serve as one functional group.

19.5 Shared and Separate Take-Up

The blackbox is taken to hold several drums, each carrying its own cord bundle, which makes each drum a separate winding channel. Cords wound on a common drum within a blackbox shorten together in a fixed ratio set by the drum radii at each cord. Cords on separate winding channels of the blackbox may shorten independently. In the language of Part II, a shared drum is a rigid constraint, equivalent to infinite coupling with a fixed offset, while separate channels leave the coordination to the oscillator coupling K_{ij} with its finite strength and its adjustable offset. Release, as defined in Chapter 8, therefore has a mechanical meaning: it requires that the released cords be on a separate channel from the retained ones.

Chapter 20

Cord and Pressure Actuation

20.1 Routing Kinematics

Let a cord pass over guides and terminate on a distal attachment, spanning joints $1, \dots, n$ with moment arms $r_1, \dots, r_n$ about the joint axes, measured from the joint axis to the line of the cord. For small motions, taking up a length s of cord produces joint displacements Δq_j constrained by $s = \sum_j r_j \Delta q_j$, and a cord tension T produces joint torques $\tau_j = r_j T$. A guide changes the cord direction without changing the cord length it contributes to that relation except through the moment arm it creates. With friction at a guide of wrap angle ϕ and friction coefficient μ, the capstan relation $\mathcal{T}_{\text{out}} = \mathcal{T}_{\text{in}} e^{-\mu\phi}$, with ϕ the total deflection angle of the cord at the guide, applies to a cord being pulled, so the tension reaching the attachment is smaller than the tension at the actuator and the loss compounds across guides.

Proposition 20.1. *Let m cords with nonnegative tensions $\mathcal{T} \in \mathbb{R}^m_{\geq 0}$ act on n joints through the structure matrix $A \in \mathbb{R}^{n \times m}$, $\tau = A\mathcal{T}$. If cords alone must produce every joint torque in $\mathbb{R}^n$, then $m \geq n+1$. This is a known requirement in the theory of tendon-driven manipulators, and a short proof is included for completeness.*

Proof. The set of achievable torques is the cone generated by the columns of A. If $m \leq n$ this cone lies in the span of at most n vectors. When the columns are independent and $m = n$, the cone is simplicial and does not contain the negative of its first generator. When they are dependent the span is a proper subspace. In either case the cone is not all of $\mathbb{R}^n$. □

Cords can only pull. A return stroke therefore requires an opposing cord, elastic recoil of the shell, or internal pressure, or some combination. With elastic or pressure return the torque relation becomes $\tau = A\mathcal{T} + \tau_{\text{el}}(q) + \tau_p(q, p)$ and the bound on m no longer applies in this form. This is the first mechanical obligation: given the guide positions and attachments, determine which joint motions a spinal pull can produce, and under what tension and shell-support conditions. It is entered in Appendix B.

20.2 Two Independent Channels

Tendon actuation and pressure actuation are separate ways of moving the body, and either can manipulate it without the other. The tendons pull on distal attachments and can produce

walking on their own. Pressure acts through the shell: when air is driven in at the center and kept flowing out through vents, the body behaves like an inflatable ornament, held up and shaped by the flow through it, and the distribution of pressure among its parts changes its posture and stiffness. The pressure layer is a full actuation channel with its own network, treated in the next chapter, and not only a return mechanism for the cords.

The two channels differ in kind. A tendon channel fixes a length and delivers a tension, so it acts like a position source. A pressure channel delivers a distributed load through the shell and sets stiffness as well as shape, so it acts like a force and compliance source. Both can carry the oscillator signals of Parts II and III. Which signals go to which channel is a design choice, and a body can use them together or in turn. The liquid-pressure variant is treated as an extension of the pneumatic channel, since the earlier writing does not specify it as the actuation system.

Chapter 21

The Pneumatic Cell Network

21.1 Cells, Gates and Reservoirs

The hexagonal cells that line the shell are treated as chambers joined by gates. A gate is a hole or diaphragm whose opening can be constricted, so that pressure moves through it differentially: the same pressure difference drives more or less flow depending on the setting. In addition, compressed-air packages are distributed through the body. Each can release its charge on an emergency trigger or after a time delay, redistributing internal pressure. These elements give the pneumatic channel three controls: the source at the center, the gate settings, and the release events.

21.2 A Linear Model

Let chamber i have pressure p_i above ambient and capacitance $C_i > 0$, let the gate between chambers i and j have conductance $g_{ij}(t) = g_{ji}(t) \geq 0$ set by its opening, let chamber i vent to ambient with conductance $v_i \geq 0$, and let $q_i(t)$ be the flow injected into chamber i from the source. Then

$$C_i \dot{p}_i = \sum_j g_{ij}(p_j - p_i) - v_i p_i + q_i. \tag{21.1}$$

This idealizes the flow through each gate as linear in the pressure difference. Real orifice flow is nonlinear and can choke, so Equation (21.1) is a model class and not a derived law.

Proposition 21.1. *Hold the gate conductances and sources constant. Let L be the weighted graph Laplacian of the gates and $V = \operatorname{diag}(v_i)$, and suppose every connected component of the gate graph contains a chamber with $v_i > 0$. Then $L + V$ is positive definite, Equation* (21.1) *converges exponentially to the unique equilibrium $p^* = (L + V)^{-1} q$, and $p^* \geq 0$ entrywise whenever $q \geq 0$.*

Proof. The quadratic form $x^\top (L + V)x = \sum_{\text{gates}} g_{ij}(x_i - x_j)^2 + \sum_i v_i x_i^2$ vanishes only if x is constant on each component and zero at a vented chamber, hence only for $x = 0$. So $L+V$ is symmetric positive definite. With $C = \operatorname{diag}(C_i)$ the matrix $C^{-1}(L+V)$ is similar to $C^{-1/2}(L + V)C^{-1/2}$, which has positive real eigenvalues, giving exponential convergence. A symmetric positive definite matrix whose off-diagonal entries are nonpositive is a nonsingular M-matrix, whose inverse has nonnegative entries, which gives $p^* \geq 0$. □

The vent condition is the mathematical form of the inflatable-ornament arrangement: air must be allowed to leave, and the flow through the body sets the steady pressure distribution. Choosing the gate conductances shapes p^{*}, which is how the gates distribute pressure differentially.

21.3 Release Events

A compressed-air package that discharges at time t_k into chamber i adds a pressure step Δ_i at that time, set by the stored mass and the chamber volume, after which Equation (21.1) relaxes the distribution at the network time constants. An emergency trigger and a time-delay trigger differ only in what fixes t_k. These events make the pneumatic channel a hybrid system, and the hybrid obligations of Appendix B apply to it as they apply to the cycling valve of Chapter 2.

21.4 Link to the Oscillator Theory

Two cycling chambers joined by a gate are coupled relaxation oscillators, and the gate conductance g_{ij} is the physical counterpart of the coupling strength K_{ij}. Constricting the gate lowers the coupling and so realizes release by decoupling. A bias in the supply to one chamber shifts its period through Equation (2.1) and realizes release by detuning. The pneumatic channel therefore carries the release operations of Part III in its own way, independent of the cords. Whether pressure alone can produce locomotion of an endomarionette is a hypothesis that has not been tested, and it is entered as an open obligation.

Chapter 22

The Segmental Spine

This chapter specifies one spinal module concretely and then the way modules join into a spine. Where the construction is undecided, the alternatives are listed and the reference choice used for the description is marked as a reference only. The concentrated thoracic arrangement of Chapter 20 is the special case of a single module.

22.1 One Module

A module is a short unit occupying the place of a vertebra or a group of vertebrae. It has five parts.

The collar is a ring or short tube of laminated paper or molded fiber whose outer edge follows the inner cross-section of the shell. In the reference arrangement the collar carries an outward flange bonded along the inside wall of the shell, continuing into the ribbing, so that loads from the cords pass into the shell through the flange. Two alternatives remain open. The collar may instead be seated in a pair of hexagonal cells and held by tabs that locate it without bonding, which keeps the shell sealed but makes the load path pass through the cell walls. Or the collar may float on flexible links with only its guides fixed to the shell, which avoids concentrating load on the shell but leaves the collar free to shift.

The guide ring is a set of smooth guide eyes mounted on the collar. Two kinds are needed. Pass-through eyes carry cords that belong to modules lower down the spine and only change their direction or hold their position as they pass. Exit eyes receive the cords that leave the module for a limb or for another part of the shell, and the first guide of each limb route is an exit eye. The guide supports carry the resultant force of Chapter 27 into the collar.

The take-up channels are the winding mechanisms of the blackbox interface. Each channel has an axle journaled in the collar and a drum or spool, or a drop reservoir in the alternative embodiment of Chapter 20, and each is driven by an actuator whose nature remains open: a motor, a pressure-driven bellows with a rack, or a ratchet with elastic return are all consistent with the interface. A channel is the unit that can be commanded separately, and Chapter 23 treats cords attached to the same channel as one constrained group.

The plenum is a chamber, one hexagonal cell or a small set of cells, fixed in the collar. It has an axial port toward each adjacent module, radial ports into the shell cells of the module's region, a vent port, a supply port from the central air line, and optionally a reservoir port with a timed or emergency release valve. Each port has an adjustable gate in the sense of

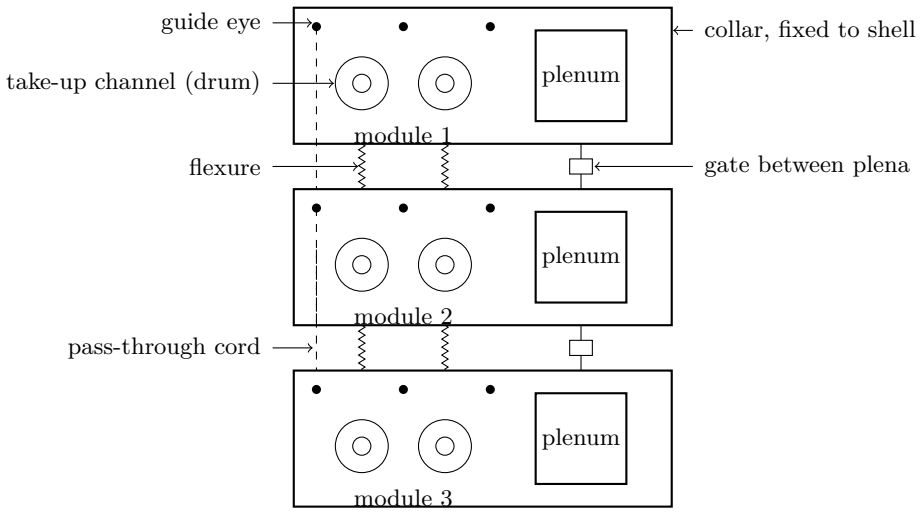


Figure 22.1: Schematic of three spinal modules. Each collar carries two take-up channels, guide eyes and a plenum. Flexures join the collars, gates join the plena, and cords belonging to lower modules pass through the eyes of the upper ones. The drawing is schematic and fixes no dimensions.

Chapter 21. The plenum is the module's contribution to the pneumatic network.

The module also carries its connections to its neighbors, described next.

22.2 Connections Between Modules

Adjacent modules are joined in three ways. The mechanical connection is a flexure between collars that carries axial load and allows bending and a limited twist. Candidate flexures are a paper bellows, a set of elastic cords, or a laminated hinge, and the choice sets how much the spine can bend and how much cord path length changes with bending. The cord connection is passive: cords of lower modules pass through the guide ring of each module above, and their path length changes with the bend of the intervening flexures, which couples spinal posture to the lengths of cords that are not attached to the spine at all. The pneumatic connection is a gated passage between the plena of adjacent modules. Its conductance is the coupling between the modules in the sense of Chapter 29, and it can be adjusted or closed.

No other connection between modules is assumed. In particular the drums of different modules are not mechanically linked, and the coordination between them is whatever the gates and the commands impose.

22.3 Forming the Spine

A spine is a stack of M modules with the flexures and gates above. In the reference arrangement the modules serve body regions as follows: the uppermost serves the head and neck, the next the shoulders and arms, the next the trunk, and the lowest the hips and legs. The number of modules and the assignment of regions are open choices. Within the hip module

each leg has its own sub-plenum with a gate to the module's common plenum, so that the legs can be coupled strongly or released from one another by closing the gate. Without this subdivision the two legs would share one chamber and could not be released independently.

22.4 A Directed Cascade Along the Spine

The snake-robot controller of Cheslet and colleagues joins its local rhythm generators by forward excitatory connections, so that each generator drives the next and the connections introduce phase differences down the body [1]. Their generators are pairs of mutually inhibitory bursting neurons, and the paper gives no phase equation for the chain. A reduced phase model inspired by that architecture, proposed here as an extension and not drawn from their paper, is the directed cascade

$$\dot\theta_1 = \omega, \qquad \dot\theta_k = \omega + k_k \sin\big(\theta_{k-1} - \theta_k - \alpha_k\big), \quad k \ge 2. \tag{22.1}$$

With the successive phase errors $e_k = \theta_k - \theta_{k-1} + \alpha_k$, the cascade gives $\dot e_2 = -k_2 \sin e_2$ and $\dot e_k = -k_k \sin e_k + k_{k-1} \sin e_{k-1}$ for $k \ge 3$.

Proposition 22.1 (Directed cascade)**.** *For positive couplings k_k the traveling pattern $e = 0$ is locally exponentially stable, because the linearization is triangular, with eigenvalues $-k_2, \dots, -k_N$. If the natural frequencies are ω_k and the leader runs at ω_1, a pattern locked to the leader exists when $|\omega_k - \omega_1| < k_k$ for every k, with*

$$\sin e_k^* = (\omega_k - \omega_1)/k_k,$$

and the branch with $\cos e_k^* > 0$ is locally exponentially stable with eigenvalues $-k_k \cos e_k^*$.

Proof. At the locked state every $\dot\theta_k$ equals the leader's rate, which gives the equations for e_k^* stage by stage. Linearizing, $\dot\delta_k = -k_k \cos e_k^* \, \delta_k + k_{k-1} \cos e_{k-1}^* \, \delta_{k-1}$, a lower-triangular system whose eigenvalues are the diagonal entries. □

Two differences from the symmetric theory should be stated. The cascade is directed: a perturbation of an early stage propagates to every later stage and none to earlier ones, whereas in the symmetric groups of Part III a disturbance spreads in both directions. The cohesion and capture theorems assume symmetric internal weights, or lower-bounded weights on a complete graph, and do not cover a directed chain, so the proposition is a separate result with a separate proof. Removing the connection into stage k frees stages k and beyond from the stages above while they remain coordinated among themselves, the directed analogue of release, and no theorem is claimed for it. A chain of pneumatic chambers joined by one-way gates is not a pure cascade, since by conservation the flow that enters a chamber leaves the one before it, as in the check-valve gate of Chapter 30.

Chapter 23

Channels, Constraints and Coupling Structure

23.1 Shared Drums and Independent Channels

A channel c has a scalar state q_c, the drum rotation. A cord i attached to channel c takes up a length $s_i = \rho_{ic} q_c$, where ρ_{ic} is the signed winding radius, and for a spool whose radius varies with the amount wound $s_i = \int \rho_{ic}(q)\,\mathrm{d}q$.

Proposition 23.1 (Cords on one drum share an actuator coordinate)**.** *All cords attached to a drum have take-up lengths determined by the single coordinate q_c, so with fixed radii the ratios of their take-up lengths are fixed. This does not fix the joint motions: the motion a cord produces depends on its routing, slack, elasticity and load, and a slack or stretched cord decouples its attachment from the drum. If q_c is driven by a periodic actuator with phase θ_c, the drum is one node of the coordination model. For two drum actuators with equal natural frequencies coupled by $K \sin(\theta_2 - \theta_1 - a)$, the phase difference tends to a for every $K > 0$, and the single-node model of a shared drum is taken as the limit $K \to \infty$; this is a definition of the limiting model. A shared drum is one actuator and not two coupled ones.*

Proof. Each take-up length is a function of q_c alone. For the coupled pair with $\Delta\omega = 0$, $\dot{\varphi} = -2K \sin(\varphi - a)$, whose stable fixed point is $\varphi = a$. □

This gives the design rule that a joint whose flexor and extensor cords are wound on opposite sides of one drum, so that one pays out as the other takes up, is driven by one actuator coordinate. Independent control of the two cords requires two drums, and whether the joint then follows the drum depends on routing, slack and elasticity. The rule also bounds the number of independent actuator coordinates of the tendon layer by the number of channels.

23.2 Which Channels Can Be Released

Two channels can be released from each other only if they are on different drums and the coupling between them is adjustable. The coupling elements available in the construction are these. A shared drum is a rigid constraint, the limiting single-node model of the previous

section and not a pair of finitely coupled actuators, and it cannot be released. A gate between the chambers that drive the two channels gives a coupling proportional to g/C with a preferred offset determined by the chamber dynamics, as derived in Chapter 29, and can be adjusted or closed. A common supply line acts as a weak global drive: changing supply pressure shifts all chamber periods at once and so changes frequencies but not couplings. The body mechanics themselves couple channels through the shell and the load, and this parasitic coupling is not modeled here and is not adjustable.

23.3 How Pressure Interacts with the Tendon Mechanism

Three interaction modes are possible and the construction has not chosen among them. In the first, pressure acts through the shell independently of the cords. The two channels then meet only in the body's mechanics, and pressure enters the theory through the output map from phases to configuration and not through the coupling. In the second, a pressure chamber drives a take-up channel, for instance through a bellows and rack, so that the channel state q_c follows the chamber volume and the oscillator phase is the phase of that chamber's cycle. The gates between such chambers are then the couplings between channels. In the third, pressure supplies return or co-contraction: a bladder at a joint opposes the cord, and for a bladder of effective area A_p acting at moment arm r_p the net torque is

$$\tau_{\text{net}} = r\mathcal{T} - r_p A_p p, \tag{23.1}$$

with the cord moment arm r and tension $\mathcal{T}$ as in Chapter 20. Raising both $\mathcal{T}$ and p stiffens the joint without moving it. The three modes enter the model at different places, the first and third in the output map and the second in the oscillator layer, The first embodiment of Chapter 24 adopts the first and third modes and keeps the second as an alternative.

23.4 The Resulting Coupling Structure

The spine induces the block structure of Part II. In the first embodiment of Chapter 24 the coupled oscillators are the pneumatic chambers and the same block structure applies to them. Channels within a module share a plenum and are strongly coupled through wide gates. Modules adjacent along the spine are coupled through the axial gates, more weakly. Modules further apart interact through a chain of such gates or through a common supply, weaker still. A body region corresponds to a group in the sense of the theorems, a set of channels whose internal coupling graph has a positive algebraic connectivity, and the cross coupling between regions is the set of gates joining their modules. The control $\mu(t)$ of Chapter 13 is the closing of those gates. Release of a group is realized by closing every gate that joins it to the rest, and this is possible in the construction only if the group's channels are not on drums shared with channels outside it.

Chapter 24

The First Embodiment

24.1 The Decision

The first worked embodiment keeps the two actuation channels independent. Drums set cord lengths, and the blackbox commands them as inputs. Gated pneumatic chambers act through the shell and provide adjustable support and return forces, through the first and third interaction modes of Chapter 23. In this embodiment the coordinated oscillators of Parts II and III are the cycling pneumatic chambers, and the drum channels are commanded set-points that are not nodes of the coupled network. Pressure-driven take-up remains an alternative embodiment, in which drum channels would become oscillator nodes coupled through gates between their driving chambers. A drum-oscillator embodiment with mechanical coordination links is not developed here.

The reason for placing the oscillators in the pneumatic layer is that its couplings are fixed by physical connections whose first-order effect has been derived in Chapter 29. The coordination of independently driven drums would instead be set by the internal mechanism of the blackbox, which is unspecified. This is a design choice made so that the theory has a definite physical architecture, and it is not a claim that the pneumatic layer is the better one. The aim of integrating a body and its controller in one pneumatic construction has a precedent in an entirely soft autonomous robot controlled by a soft microfluidic oscillator [23], which uses materials and manufacturing methods very different from paper shells and cords. This monograph does not claim electronics-free control, since the blackbox is left unspecified.

24.2 What Each Gate Changes

Each chamber cycles as in Chapter 2 and is connected to a bladder at a joint. The elements and what opening or closing them changes are as follows.

Element	Opening changes	Model quantity
Coordination gate between chambers i, j	Coupling strength, in proportion to g_{ij}/C for small g_{ij}	K_{ij}
Support gate, chamber to bladder	Gain and bandwidth of the force transmitted to the joint; loads the chamber and shifts its frequency	τ_p, detuning
Return gate, bladder vent	Return rate and the pressure the bladder holds	Return stiffness
Supply branch gate	Supply pressure p_s at a group	Frequency, by detuning
Reservoir valve	An impulsive pressure step in a chamber	Phase kick

For a bladder of capacitance C_b with support gate conductance g_b and vent conductance v_b, the bladder pressure obeys $C_b\dot{p}_b = g_b(p - p_b) - v_b p_b$, so its steady gain from the chamber is $g_b/(g_b + v_b)$ and its cutoff rate is $(g_b + v_b)/C_b$. The joint torque is $\tau_p = r_p A_p p_b$. The supply pressure sets the chamber frequency through the period of Equation (2.1): differentiating gives

$$\frac{\mathrm{d}\omega}{\mathrm{d}p_s} = \frac{2\pi}{T^2}\,\tau_c \left[\frac{1}{p_s - p_h} - \frac{1}{p_s - p_l}\right] > 0, \tag{24.1}$$

so a higher supply pressure raises the frequency. A reservoir discharge of size δp_i into chamber i shifts its phase to first order by $Z(\theta_i)\delta p_i$ in time units, so the angular shift is at most $\zeta = (2\pi/T) Z_{\max} \delta p_{\max}$, with Z the response curve of Chapter 29. An impulsive kick is an event for Proposition 18.1, and for a state with spread at most γ_{in} the kick keeps the spread within γ_a whenever $\gamma_{\text{in}} + 2\zeta \le \gamma_a$. This is a direct sufficient condition and is less restrictive than the coherence-based condition of Part IV.

24.3 Residual Coupling

Closing a gate establishes pneumatic isolation only to the extent that the other paths are negligible or bounded. Three paths remain.

Gate leakage means a closed gate keeps a minimum conductance, and since the coupling is proportional to conductance a residual coupling $K^{\text{res}}_{ij} \ge 0$ remains. The common supply couples chambers through the line impedance: flow drawn by one chamber lowers the supply pressure of others, so each chamber sees $p_s + \delta_i(t)$ with $|\delta_i| \le \bar{\delta}$, which by the derivative above perturbs its frequency by at most $L_s\bar{\delta}$, where L_s bounds $|\mathrm{d}\omega/\mathrm{d}p_s|$ over the operating range. The body couples the layers through the mechanics: a tendon-driven joint motion changes the bladder volume, and for an isothermal gas the pressure rate perturbation is $-(p/V)\dot{V}_b$, which perturbs the phase velocity by at most $\zeta_b = (2\pi/T) Z_{\max} (p_{\max}/V_{\min}) \dot{V}_{b,\max}$. The supply and body paths are modeled as bounded frequency perturbations $d_i(t)$ with $|d_i| \le L_s\bar{\delta} + \zeta_b$, and the leakage as a floor on the cross weights, so that $K_{ij}(t) = K^{\text{res}}_{ij} + \mu(t)\bar{K}^{\text{adj}}_{ij}$. Let ς_a bound the group mean of $|d_i|$ over group a.

Proposition 24.1 (Release and recapture with residual paths)**.** *In the setting of Theorem 13.3, let the cross degrees $\bar{d}^a$ include the residual weights, let $\Phi_a = \Omega_a + \|\bar{d}^a\|_2 + \|d\|_{2,a}$ with $\|d\|_{2,a}$ the Euclidean bound on the perturbations d_i, and define γ_a as before. Let N_{res} and*

$\bar{N}_{\mathrm{adj}}$ *be the effective couplings of the residual and adjustable weights, and* $\bar{N} = N_{\mathrm{res}} + \bar{N}_{\mathrm{adj}}$. *Then* $D_a(t) \le \gamma_a$ *for every control. Release is guaranteed on an interval where* $\mu = \mu_{\mathrm{rel}}$ *provided*

$$\big(N_{\mathrm{res}} + \mu_{\mathrm{rel}}\bar{N}_{\mathrm{adj}}\big)(1 + \gamma_g + \gamma_h) < |\Delta\bar{\omega}| - \varsigma_g - \varsigma_h, \tag{24.2}$$

and recapture with $\mu = 1$ *is guaranteed in the sense of Theorem 9.1 provided* $s^* = (|\Delta\bar{\omega}| + \varsigma_g + \varsigma_h)/\bar{N} + \gamma_g + \gamma_h < 1$.

Proof. The group mean frequencies are perturbed by the group means of d_i, so the effective mismatch lies within $\varsigma_g + \varsigma_h$ of $\Delta\bar{\omega}$. The reduced equation is $\dot{u} = \Delta_{\mathrm{eff}} - N(t)\sin u + r$ with $N(t) = N_{\mathrm{res}} + \mu\bar{N}_{\mathrm{adj}}$ and $|r| \le N(t)(\gamma_g + \gamma_h)$. In release, $|\dot{u}| \ge |\Delta\bar{\omega}| - \varsigma_g - \varsigma_h - N(1 + \gamma_g + \gamma_h)$, which is positive under the stated condition. In recapture, $\dot{u} \le |\Delta\bar{\omega}| + \varsigma_g + \varsigma_h + N(\gamma_g + \gamma_h) - N\sin u$, and the capture argument applies with the stated s^*. □

The proposition shows what isolation costs. Closing every adjustable gate guarantees release only if the residual coupling and the residual frequency perturbations together are smaller than the frequency mismatch between the groups, and when that fails a detuning drive from the supply branches must supply the margin. The three-group example below inherits this requirement.

Chapter 25

Worked Example: Stance, Free Leg and Arms

25.1 The Three Groups

Take three groups of channels: s, serving the trunk and one leg that is intended to carry the body; f, serving the other leg; and a, serving the arms. In the reference spine s and f share the hip module, with the sub-plena of the previous chapter, and a is served by the shoulder module. The groups have internal coupling graphs with algebraic connectivities $\lambda_s, \lambda_f, \lambda_a$ and pairwise cross weights $\bar{K}^{(sf)}, \bar{K}^{(sa)}, \bar{K}^{(fa)}$ scaled by controls $\mu_{sf}(t), \mu_{sa}(t), \mu_{fa}(t) \in [0,1]$ that open and close the gates. The cross coupling is reciprocal and the preferred offsets are consistent as in Chapter 13. Let $\bar{d}^a$ be the vector of total cross coupling of each member of group a to all other groups, and define $\Phi_a = \Omega_a + \|\bar{d}^a\|_2$ and $\gamma_a = \arcsin(\sqrt{2}\Phi_a/\lambda_a)$ as before, assuming $\sqrt{2}\Phi_a < \lambda_a$.

25.2 Preservation of Internal Coordination

Theorem 25.1 (Three-group cohesion and reduction)**.** *If $\|e_a(0)\|_2 \le \gamma_a/\sqrt{2}$ for each group, then for every choice of the controls $\mu_{ab}(t) \in [0,1]$ the spreads satisfy $D_a(t) \le \gamma_a$ for all t. The collective phases satisfy $\dot{\bar{\vartheta}}_a = \bar{f}_a + \sum_{b \neq a}(S_{ab}(t)/n_a)\sin(\bar{\vartheta}_b - \bar{\vartheta}_a - a_{ab}) + \rho_a$ with $S_{ab}(t) = \mu_{ab}(t)\sum \bar{K}_{ij}$ and $|\rho_a| \le \sum_{b \neq a}(S_{ab}/n_a)(D_a + D_b)$. In particular, if every control involving group a is zero and no residual cross coupling or other disturbance acts on a, then $\rho_a = 0$ and $\dot{\bar{\vartheta}}_a = \bar{f}_a$ exactly.*

Proof. The external perturbation on member $i \in a$ is bounded by its total cross degree, so Theorem 13.1 applies to each group with the degree vector $\bar{d}^a$, for every control. The collective phase equation follows by averaging over the group, with the internal terms cancelling by symmetry, and the remainder is bounded as in Theorem 13.2. If all controls involving a vanish and no residual path acts on a, the cross terms are absent. □

This is the preservation result, and it concerns internal coordination only: whichever group is released, the others keep their spreads within the stated bounds. A released group advances at its own mean frequency exactly when no residual path acts on it, which in the first embodiment is never strictly true, and otherwise within the bounds of Proposition 24.1.

25.3 Release and Recapture on a Pairwise Schedule

Recapture with all three links active at once is treated in Section 25.4, under consistent targets and a margin condition. A schedule that keeps at most one pair of groups coupled at any time reduces the problem to the two-group theorem, which needs neither.

Theorem 25.2 (Pairwise-exclusive schedule)**.** *Suppose at every time at most one control μ_{ab} is nonzero. On an interval where the active pair is (a, b), the relative phase $u = \bar{\vartheta}_b - \bar{\vartheta}_a - a_{ab}$ obeys $\dot{u} = \Delta\bar{\omega}_{ab} - N_{ab}(t)\sin u + r$ with $N_{ab} = \mu_{ab}\bar{S}_{ab}(1/n_a + 1/n_b)$ and $|r| \le N_{ab}(\gamma_a + \gamma_b)$, and the third group advances at its own mean frequency, up to the residual paths of Proposition 24.1 when they are present. Hence the release and recapture statements of Theorem 13.3 hold for the active pair on that interval, with $s^*_{ab} = |\Delta\bar{\omega}_{ab}|/\bar{N}_{ab} + \gamma_a + \gamma_b < 1$ and with the restoration time taken at the start of the interval.*

Proof. On the interval the third group is isolated, so, absent residual paths, its collective phase advances at its mean frequency by Theorem 25.1. The active pair then obeys the two-group reduction of Theorem 13.2 with spreads bounded by γ_a, γ_b, and the release and capture arguments apply to it. □

The theorem is conservative, since γ_a is computed with the full cross degree to both other groups, although only one pair is active at a time. It describes the sequence needed in the example: the free leg is coupled to the stance group while the arms are released, then released from the stance group, then recaptured, and then the arms are coupled to the stance group while the free leg is released.

25.4 Simultaneous Recapture

When all three links operate at once the pairwise theorem cannot simply be applied twice, because each pair is acted on by the third group. Write the targets as group offsets B_s, B_f, B_a and set $a_{ab} = B_b - B_a$, so that the pairwise offsets are consistent around the cycle: $a_{sf} + a_{fa} + a_{as} = 0$ modulo 2π. Consistency is automatic when the offsets come from group targets, and it is a real condition when the offsets are imposed pair by pair by physical connections. With $u_{sf} = \bar{\vartheta}_f - \bar{\vartheta}_s - a_{sf}$ and $u_{fa} = \bar{\vartheta}_a - \bar{\vartheta}_f - a_{fa}$, and $u_{sa} = u_{sf} + u_{fa}$ by consistency, the reduced equations of Theorem 25.1 give, writing $W_{ab} = S_{ab}/n_a$ and $\Delta_{ab} = \bar{f}_b - \bar{f}_a$,

$$\dot{u}_{sf} = \Delta_{sf} - (W_{sf} + W_{fs})\sin u_{sf} + W_{fa}\sin u_{fa} - W_{sa}\sin u_{sa} + \rho_f - \rho_s, \tag{25.1}$$

$$\dot{u}_{fa} = \Delta_{fa} - (W_{fa} + W_{af})\sin u_{fa} + W_{fs}\sin u_{sf} - W_{as}\sin u_{sa} + \rho_a - \rho_f. \tag{25.2}$$

The terms in $\sin u_{fa}$ and $\sin u_{sa}$ in the first equation are the influence of the third group on the pair (s, f), and similarly for the second equation. They are not small compared with the diagonal terms in general, which is why the problem is treated as a three-node network.

Treat the shifted collective phases $\tilde{\vartheta}_c = \bar{\vartheta}_c - B_c$ as three effective oscillators with coupling weights W_{ab}, which are not symmetric.

Theorem 25.3 (Simultaneous recapture of three groups)**.** *Assume consistent targets, all three links operating with every $S_{ab} > 0$, and the spread bounds $D_c \le \gamma_c$ of Theorem 25.1. Let $\kappa_{\min} = \min_{a\neq b} S_{ab}/n_a$, let W_f be the spread of the group mean effective frequencies, and*

let $R = \max_a\big[\sum_{b\neq a}(S_{ab}/n_a)(\gamma_a+\gamma_b)+\varsigma_a\big]$*, with* $\varsigma_a = 0$ *when no residual perturbations are present. If* $3\kappa_{\min} > W_f + 2R$*, set* $\gamma^G = \arcsin\big((W_f+2R)/(3\kappa_{\min})\big)$ *and let* $D^G(t)$ *be the spread of* $\tilde\vartheta_s, \tilde\vartheta_f, \tilde\vartheta_a$*. If* $D^G(t_r) \le \gamma^G$ *at the restoration time then* $D^G(t) \le \gamma^G$ *afterward; if* $D^G(t_r) < \pi - \gamma^G$ *then* $\limsup D^G \le \gamma^G$*, and for every* $\eta > 0$ *with* $\gamma^G + \eta < D^G(t_r)$ *the spread enters and stays in* $[0, \gamma^G + \eta]$ *after a finite time depending on* $D^G(t_r)$ *and* η*. In particular every pairwise relative phase satisfies* $|u_{ab}| \le D^G$*, and every cross pair of oscillators satisfies* $\limsup|\vartheta_j - \vartheta_i - a_{ab}| \le \gamma^G + \gamma_a + \gamma_b$ *modulo* 2π*.*

Proof. The reduced equations read $\dot{\tilde\vartheta}_a = \bar f_a + \sum_{b\neq a} W_{ab}\sin(\tilde\vartheta_b - \tilde\vartheta_a) + \rho_a$ with $|\rho_a|$ at most R once the residual perturbations are included. Apply Proposition 13.1 in its time-dependent form to these three nodes: the extreme nodes satisfy $D^+D^G \le (W_f+2R) - \kappa_{\min}\cdot 3\sin D^G$ for $D^G \le \pi$, because each weight is at least $\kappa_{\min}$ and the sines are nonnegative. The invariance and attraction conclusions follow as in Theorem 6.1. For the last statement combine $|u_{ab}| = |\tilde\vartheta_b - \tilde\vartheta_a| \le D^G$ with the within-group deviations. □

The theorem requires the minimum weight to dominate, so it is not implied by the pairwise theorem, and it does not imply it: the pairwise condition compares each mismatch with its own pair's coupling, and this one compares the whole spread of mean frequencies with three times the weakest weight. When one link is much weaker than the others the pairwise schedule is the better tool. If the preferred offsets imposed by the physical connections are not consistent around the cycle, no common target exists, the network is frustrated, and the theorem does not apply; a bound for the frustrated case is open. In the first embodiment, purely diffusive gates between identical chambers prefer a zero offset, so any consistent assignment other than all-in-phase requires gates with designed delays or crossed connections whose offsets must be derived.

25.5 A Simulation

Groups of six, five and five oscillators with sparse internal rings, internal weights scaled to give $\lambda_s = 8.90$, $\lambda_f = 14.08$ and $\lambda_a = 13.41$, and cross weights up to 0.18 were simulated. The bounds were $\gamma_s = 0.192$, $\gamma_f = 0.160$, $\gamma_a = 0.132$, and the pairwise conditions held with $\gamma^*_{sf} = 0.566$ and $\gamma^*_{sa} = 0.522$. The schedule coupled s and f for 40 time units, released everything for 50, recoupled s and f for 50 after setting their relative phase to 1.5, and then coupled s and a for 50. The largest observed spreads were 0.078, 0.069 and 0.056. At the end of the third stage the relative phase of f to s was 0.186, and at the end of the fourth stage that of a to s was -0.174, both inside the corresponding arcs. The runs check the theorem at one parameter set and do not test its boundaries.

For simultaneous recapture the same groups were coupled through all three links at once. The weakest weight was $\kappa_{\min} = 0.181$, the spread of mean frequencies $W_f = 0.162$, and $R = 0.175$, so that $(W_f + 2R)/(3\kappa_{\min}) = 0.942$ and $\gamma^G = 1.229$, which leaves a guaranteed capture region of $D^G < 1.913$. Started with a group-level spread of 1.7, inside the guarantee, the late spread was 0.208. Three further starts with spreads of 2.9, 4.0 and 3.0, outside the guarantee, also converged on the circle to the same relative phases, 0.111 and -0.097, the start with spread 4.0 after a full phase slip. These outside runs show that the guarantee is conservative at this parameter set and say nothing about other parameter sets.

25.6 What the Theorem Does Not Establish

Calling a group stance does not establish that it supports the body. The theorems concern phase coordination, and balance additionally requires mechanical and contact conditions that they do not address: the position of the center of mass relative to the support, the contact forces and friction at the foot, the torque limits of the actuators, and the inertial effects of the moving free leg and arms. These can be linked to the theorem only through a stated hypothesis. Let Π map the phases to a body configuration, and let $\mathcal{S}$ be the set of configurations satisfying the chosen static balance and contact conditions. The hypothesis of stance admissibility is that Π maps every phase state compatible with the spread bounds of the retained groups, together with every phase of the released groups, into $\mathcal{S}$. It is a requirement on the body's mechanics and is not a consequence of the oscillator theory, and for one-leg standing with moving arms it is demanding. Dynamic balance requires the equations of motion and contact model and lies outside this monograph.

Chapter 26

Fabrication Routes

Two soft-robotics precedents bear on this chapter. Soft valves with a bistable membrane are fabricated by molding [15], which suggests a manufacturing route for valve elements, and an integrated design and fabrication strategy for an entirely soft autonomous robot is reported in [23]. Both use elastomers, and neither establishes that a paper-shell body or its valves can be made. The material available for fabrication is a set of manufacturing concepts. None has been carried out, and the dimensional and procedural details are unspecified: shell thickness, guide mounting, cord routing at particular joints, sealing and assembly order. This chapter therefore presents alternative routes for each part of the body and states what each would have to deliver, and it does not describe them as one recipe.

26.1 Shell Forms

The shell is a hollow body of paper, and three forms are considered. A folded shell is made from sheets creased along designed lines so that the surface develops into a closed or partly closed form, which suits joints because the creases can act as hinges. A laminated shell is built from several layers bonded together, which allows the layers to differ, for instance in grain direction or in whether a layer carries embedded cord. A paper-mâché shell is built up from strips and paste over a form or a sacrificial core, which suits complex curved surfaces and can be made in one piece. These can be combined within a body, for example folded joints in laminated or molded limbs.

Reinforcement takes two forms in the notes. Ribbing is a set of stiffer members attached to or formed into the shell, running along the directions of the principal loads from the cords. Hexagonal cells are a honeycomb arrangement whose walls both stiffen the shell and serve as the chambers of Chapter 21, which makes the reinforcement structure and the pneumatic network the same object. This coincidence is a design hypothesis: it requires that the cell walls be both strong enough to carry shell loads and sealed well enough to hold pressure, and the two requirements may pull in different directions.

26.2 Four Forming Concepts

Four processes appear in the material. They are alternatives, and some may combine.

Pressure-assisted cellulose forming presses or inflates wet cellulose fiber against a mold or within a constraint, so that fibers consolidate into a shell whose shape is set by the constraint and whose density is set by the applied pressure. The unspecified choices are the pressure, the drying procedure and the mold material.

Intermittent steam inflation forms the shell by inflating it with steam in repeated intervals, with cable reinforcement laid in before or between intervals. The idea is that steam softens the fibers and bonding agent so the shell can be shaped from inside, then sets as it cools and dries. The unspecified choices are the interval schedule, the temperature and the way cables are held in position.

Magnet-directed reinforcement places reinforcing elements by using magnets to position or orient them within a wet layer, so that reinforcement lies along intended load paths without hand placement. This requires reinforcing elements that respond to magnets, and the compatibility of such elements with paper and with the later pneumatic function is unspecified.

A sacrificial honeycomb scaffold is a honeycomb that is proposed to soften or dissolve into an inner layer, leaving a smooth lining or a cellular interior depending on how it is treated. This connects the hexagonal cells of the body to fabrication, since the scaffold would define the cells before becoming part of the shell. The material, the dissolving agent and the resulting lining properties are unspecified.

26.3 What Fabrication Must Deliver

Whatever route is used, it must deliver three things that can be stated as inequalities, using standard thin-wall mechanics. For a thin-walled cylindrical chamber of radius r and wall thickness $t \ll r$ under internal gauge pressure p, the hoop stress is $\sigma = pr/t$, and for a thin spherical chamber it is $\sigma = pr/(2t)$. A shell material with allowable tensile stress σ_{allow} therefore limits the pressure to

$$p \le p_{\max} = \frac{\sigma_{\mathrm{allow}}\, t}{r} \quad \text{(cylinder)}, \tag{26.1}$$

with a factor of two more for a sphere. The thresholds and supply pressure of the oscillators of Chapter 2 must satisfy $p_s \le p_{\max}$ for every chamber in which they act, with a safety margin. The value of σ_{allow} for the laminate, including the effect of seams and any seal layer, is unspecified and must be measured.

Chapter 27

Loads on the Shell

27.1 Guides

A cord of tension $\mathcal{T}$ that changes direction by a deflection angle ϕ at a guide presses on the guide with a resultant force of magnitude $2\mathcal{T}\sin(\phi/2)$ when friction is neglected. The guide supports, which are fixed to the shell, must carry this force into the shell. If the shell material and mounting have an allowable bearing stress σ_{bear} over a mounting patch of area A, the requirement is

$$2\mathcal{T}\sin(\phi/2) \le \sigma_{\text{bear}} A. \tag{27.1}$$

Friction at a smooth sliding eye changes the tension across the guide by the factor $e^{-\mu\phi}$, which means the tension at the attachment is smaller than at the blackbox, and the load on later guides is smaller. The resultant must then be computed from the two unequal tensions as a vector sum.

27.2 Attachments

An attachment transmits the cord tension $\mathcal{T}$ into the shell. It must spread the load over an area A_{att} large enough that $\mathcal{T} \le \sigma_{\text{bear}} A_{\text{att}}$, and it must do so without tearing or peeling the laminate. Because an attachment terminates the cord at a point on a distal part of the limb, its position fixes the moment arm in the routing kinematics of Chapter 20, and the placement of the attachment is therefore a kinematic as well as a structural choice.

27.3 Shell Stability

A thin shell under cord loads and internal pressure can fail by buckling or by local collapse near guides and attachments, and a pressurized shell can fail by hoop stress as above. The conditions under which the shell keeps its form under the combined loads of tendon tension and internal pressure have not been derived here. They are the mechanical feasibility obligation of Appendix B.

Chapter 28

Related Mechanisms and Precedents

This chapter places the construction beside three families of existing mechanisms: paper and origami structures, tendon transmission, and pneumatic muscles. Each is a precedent for a part of the design and none establishes the design as a whole. The descriptions of the cited works in this chapter rest on their published abstracts and bibliographic records; the full texts have not been checked against these statements, and the figures quoted from them, such as the force expression of the McKibben model, are taken from the literature and not reproduced here.

28.1 Paper and Origami Structures

Rus and Tolley review origami robots, in which folding is combined with smart-material actuators so that actuation is intrinsic to the folded structure, which decouples design from fabrication complexity; the review covers design principles, fabrication, actuation, materials and control [16]. Felton and colleagues built a crawling robot that starts as a flat composite sheet with embedded electronics and shape-memory hinges and folds itself into a functional machine [9]. That work supports autonomous folded assembly from a flat sheet. It does not establish pneumatic paper-shell construction, since its folding is driven by heated hinges and not by pressure. Martinez and colleagues make pneumatic actuators from composites of paper and an elastomer, in which the paper constrains the elastomer so that inflation produces anisotropic motion, bending, extension, contraction or twisting, and folded paper increases stiffness and anisotropy [12]. These are precedents for folding as a fabrication route, for paper as a structural and constraining layer, and for pressure as an actuation mechanism acting through a paper-reinforced body, which are the ideas behind the folded and laminated shell forms of Chapter 26. Two differences matter. The pneumatic actuators of Martinez and colleagues obtain a sealed chamber from the elastomer, whereas a paper shell has no seal of its own, so sealing is an open choice in the design table and not a solved one. And none of these works uses cords drawn from a central spine or cell walls that serve as both reinforcement and chambers.

28.2 Tendon Transmission

Tendon-driven manipulators have a structural theory in which the tendon routing is described by a structure matrix. Lee and Tsai define the structural isomorphism of tendon-driven manipulators, describe the structural characteristics of such mechanisms, and use them to enumerate mechanisms having a pseudotriangular structure matrix with up to six degrees of freedom [10]. That work is cited here for the structure-matrix formalism and not for the cord count. The requirement of Proposition 20.1, that cords which can only pull need at least $n + 1$ cords for full control of n joints, is a standard result of the tendon-driven manipulator literature, and the proof here is included only for completeness. A standard fact from convex geometry sharpens it: the columns of the structure matrix positively span the torque space exactly when they span it and some strictly positive combination of them is zero, that is, when the structure matrix has rank n and a null vector with all entries positive. This monograph proves only the count. What is different here is the setting: cords pass through smooth guide eyes that slide and do not roll, the drums are grouped into channels that share or do not share actuator coordinates, and the same body carries a pressure channel.

28.3 Pneumatic Muscles

The McKibben muscle is a pressurized bladder constrained by a braided sleeve, which contracts and develops tension when the bladder is inflated, and Chou and Hannaford report mechanical testing and modeling of it, including an idealized static model [8]. Their measurements show hysteresis between tension and length, so the muscle is not a simple pressure-to-length converter. In that model the force is $F = (\pi D_0^2 P'/4)(3\cos^2\vartheta - 1)$, with P' the gauge pressure, ϑ the braid angle and $D_0 = b/(n\pi)$ the diameter at $\vartheta = 90°$, so the force falls to zero when $\cos^2\vartheta = 1/3$; the model neglects friction and end effects, which the measured hysteresis reflects. A McKibben muscle is an alternative mechanism for the pressure-driven take-up of the second interaction mode of Chapter 23: it would replace the drum and the cord of a channel by a contractile element whose force is set by the chamber pressure and its contraction. Using it would require a model that includes the hysteresis and not only a static pressure-to-length map. It does not enter the first embodiment, in which drums set cord lengths independently of pressure, and no coupling law or interface error budget is derived for it here. The paper-elastomer actuators of the previous section are a second alternative, for a pressure channel that acts through the shell.

Mechanism	What sets the motion	Where it would enter	Not established here
Drum and cord, as in the first embodiment	Drum coordinate through routing	Interface chapter	Joint motion
McKibben muscle	Pressure and contraction	Pressure-driven take-up, mode two	Coupling and interface
Paper-elastomer composite	Pressure through anisotropic shell	Pressure channel through the shell	Seals in an all-paper body

Chapter 29

From Chamber Coupling to Phase Coupling

The oscillator theory assumes sinusoidal coupling with a strength K_{ij} and a preferred offset α_{ij}. For the pneumatic oscillator these quantities must be derived from the chamber model and not assigned by analogy. This chapter gives the first-order derivation for two chambers joined by a gate, with its limitations. A systematic phase reduction for limit cycles in hybrid systems, which combine continuous evolution with discrete switching and carry the effect of a perturbation across a switching event, is developed by Shirasaka and colleagues [18]. The response curve derived below is obtained directly for the piecewise-linear cycle, where the pressure is continuous at the switching events and no jump term arises. For a chamber with a single pressure variable the phase equation obtained from it is exact (Theorem 34.1), so the general construction of Shirasaka and colleagues is not needed for the model of this Part; it would be needed for chambers with more state.

29.1 Phase Response of the Pressure Cycle

Use the relaxation cycle of Chapter 2, with phase $\theta \in [0, T)$ measured in time from the moment the valve closes, so that charging occupies $[0, t_c]$ and discharge $[t_c, T]$, where $t_c = \tau_c \ln \frac{p_s - p_l}{p_s - p_h}$. A small pressure perturbation δp applied at phase θ decays with the time constant of the current phase until the next switching event, and shifts that event by the perturbation divided by the pressure rate at the event. This gives the phase response to a pressure perturbation, positive for an advance:

$$Z(\theta) = \begin{cases} \dfrac{e^{-(t_c-\theta)/\tau_c}}{\dot{p}_c}, & 0 \le \theta < t_c, \\ -\dfrac{e^{-(T-\theta)/\tau_o}}{|\dot{p}_l|}, & t_c \le \theta < T, \end{cases} \qquad \dot{p}_c = \frac{p_s - p_h}{\tau_c}, \quad |\dot{p}_l| = \frac{p_l - p_o}{\tau_o}. \tag{29.1}$$

A rise in pressure during charging advances the cycle, and a rise during discharge delays it.

29.2 Weak Coupling Through a Gate

Let two identical chambers of capacitance C be joined by a gate of small conductance g, so that chamber i receives an inflow $g(p_j - p_i)$. To first order in g/C the phases obey $\dot{\theta}_i = 1 + (g/C) Z(\theta_i)\big[p(\theta_j) - p(\theta_i)\big]$, where $p(\theta)$ is the unperturbed pressure waveform. Averaging over a cycle, the phase difference $\varphi = \theta_2 - \theta_1$ obeys

$$\dot{\varphi} = \frac{g}{C}\big[\bar{\Gamma}(-\varphi) - \bar{\Gamma}(\varphi)\big], \qquad \bar{\Gamma}(\psi) = \frac{1}{T}\int_0^T Z(s)\big[p(s+\psi) - p(s)\big]\, ds. \tag{29.2}$$

The coupling strength is proportional to g/C, which is the relation between a physical parameter and the coupling in the oscillator theory. The coupling function is not a sinusoid. A sinusoidal approximation keeps its first Fourier harmonic, whose amplitude gives an effective K and whose phase gives an effective α, and the theorems of Part III then apply to that approximation, but the higher harmonics are neglected and their effect is not bounded. This is an open step.

For the cycle with $p_s = 10$, $p_h = 6$, $p_l = 3$, $p_o = 0$, $\tau_c = 1$, $\tau_o = 0.5$ (arbitrary pressure units), evaluating Equation (29.2) numerically gives $\dot{\varphi} < 0$ for $0 < \varphi < T/2$ and $\dot{\varphi} > 0$ for $T/2 < \varphi < T$, so the in-phase state is stable and the antiphase state is unstable. A direct simulation of two coupled chambers with $g/C = 0.05$, started at relative phases $0.05T$, $0.3T$ and $0.5T$, converged to in-phase in all three runs, which agrees with the prediction for this parameter set. The conclusion is specific to that set and to gate coupling that is purely diffusive in pressure.

29.3 Consequences for Design

Purely diffusive gate coupling between identical chambers therefore realizes a preferred offset near zero, so gaits requiring other offsets, such as antiphase leg motion, need a different arrangement. Chapter 30 derives the interaction function for three connection mechanisms, tests them against hybrid simulations, and states which offsets each can realize.

Chapter 30

Interaction Functions and Gate Offsets

30.1 The Interaction Function

Take two chambers with the cycle of Chapter 2, phases $\theta_1 = s$ and $\theta_2 = s + \varphi$ measured in time from valve closing, and a connection that delivers inflows q_1, q_2 to the chambers, proportional to a conductance g. Write $\epsilon = g/C$ and let $\delta = \omega_2/\omega_1 - 1$ be the relative natural rate difference of the chambers. To first order in ϵ, averaging over a cycle, the phase difference obeys

$$\dot{\varphi} = \delta + \epsilon\, G(\varphi), \tag{30.1}$$

where G is the cycle average of $Z(\theta_2)\hat{q}_2 - Z(\theta_1)\hat{q}_1$ with $\hat{q}_i = q_i/g$, and Z is the phase response of Chapter 29. In the notation of Part II the interaction of chamber i with chamber j is $H(\psi)$ with $\psi = \theta_j - \theta_i$, so that $G(\varphi) = H(-\varphi) - H(\varphi)$ for a symmetric connection of identical chambers. Only the odd part of H enters, and the shifted sinusoid $K\sin(\varphi - \alpha)$ of the oscillator theory is an approximation to $-\frac{1}{2}G$ in angle variables whose accuracy is checked below.

Two facts hold for all cycle parameters. The first is a normalization: for the time-phase variable $\dot{\theta} = 1$ along the cycle, so $Z(\theta)\dot{p}(\theta) = 1$, which can be checked directly on both branches of the cycle. The second concerns symmetry.

Proposition 30.1 (Symmetric connections)**.** *For two identical chambers joined by a symmetric gate, with or without a transport delay in both directions, G is odd, so $\varphi = 0$ and $\varphi = T/2$ are equilibria of* (30.1) *at $\delta = 0$ for every delay. Symmetry does not exclude further zeros of G, and whether further equilibria exist has to be read from the derived function.*

Proof. Exchanging the roles of the two chambers sends φ to $-\varphi$ and G to $-G$, because each chamber receives the other's pressure through the same gate and delay. Oddness together with periodicity gives $G(T/2) = -G(-T/2) = -G(T/2) = 0$. □

For the cycle used below, evaluated on a grid of step $T/480$ in φ, the diffusive gate has no zeros of G other than 0 and $T/2$. For the delayed gate, scanned over τ/T in steps of 0.01, no further zeros were found for $0.34 \le \tau/T \le 0.80$, the window in which antiphase is stable and in-phase unstable, while further zeros appear for τ/T between 0.31 and 0.33 and between

0.81 and 0.83, near the edges of that window, where their offsets move continuously with the delay. Their stability was not characterized. The statement that a symmetric connection realizes only in-phase or antiphase therefore holds for the diffusive gate and for the delayed gate inside the window, as a property of these computed functions and not as a consequence of symmetry.

A delay shifts the argument of the interaction function by the delay in time units, which gives $G_\tau(\varphi) = \bar{\Gamma}(-\varphi-\tau) - \bar{\Gamma}(\varphi-\tau)$ with $\bar{\Gamma}(\psi) = \langle Z(s)[p(s+\psi)-p(s)]\rangle$. A delay is therefore a fixed time, not a fixed phase: the phase it represents is $2\pi\tau/T$, and it changes when the period changes.

All three mechanisms below were evaluated numerically for the cycle $p_s = 10$, $p_h = 6$, $p_l = 3$, $p_o = 0$, $\tau_c = 1$, $\tau_o = 0.5$ (arbitrary units, $T = 0.906$) and compared with a hybrid simulation of two chambers with $\epsilon = 0.08$, in which the valves switch at the thresholds and the connecting flow acts continuously.

30.2 An In-Phase Connection

The first mechanism is the symmetric diffusive gate, $q_1 = g(p_2 - p_1)$ and $q_2 = g(p_1 - p_2)$.

Proposition 30.2 (Local strength of the diffusive gate)**.** *For two identical chambers joined by a symmetric diffusive gate, with the phase difference φ measured in the time unit of the cycle and $\epsilon = g/C$ in the inverse of that unit, the interaction function of the first-order averaged equation* (30.1) *satisfies $G'(0) = -2$, so near the in-phase state $\dot{\varphi} = -2\epsilon\varphi$ to first order, for every cycle. The scalings are stated in Section 37.3.*

Proof. $G(\varphi) = \langle (Z(s) + Z(s+\varphi))(p(s) - p(s+\varphi))\rangle$, so $G'(0) = -2\langle Z\dot{p}\rangle = -2$ by the normalization $Z\dot{p} = 1$. □

In the Kuramoto form $\dot{\varphi} = -2K\sin\varphi$ the local strength is therefore $K_{\mathrm{loc}} = g/C$ in the first-order model. The conductance over capacitance sets the scale of the coupling rate, and here it also equals the local strength, but this is a property of the slope at the in-phase state and does not describe the whole interaction function. The Fourier expansion of G in the angle $x = 2\pi\varphi/T$ has sine coefficients -0.2098, -0.0144, -0.0021, -0.0042, -0.0002 for harmonics one to five, so the first harmonic fit $G \approx -0.2098\sin x$ has a peak error of 8.5 percent of the peak of G, and its effective strength is $K_1 = 0.727\,g/C$. The first harmonic therefore underestimates the local slope: it gives 1.455 where the true value is 2, an error of 27 percent, while the peak of G, 0.2084, agrees with the first-harmonic amplitude 0.2098 to within one percent. The sinusoid is accurate for the locking range, $|\delta| \le 0.208\,\epsilon$, and inaccurate for the relaxation rate near lock. In the hybrid simulation with $\epsilon = 0.08$ a phase difference started at $0.06T$ decayed at the rate 0.1544 (Chapter 37). That is 3.5 percent below the tangent prediction $2\epsilon = 0.160$. The first-order model fitted over the same window gives 0.1539, which supports attributing the deficit of that fit to the finite amplitude of the displacement, since $G(\varphi)/\varphi$ is -1.73 at $\varphi = 0.06T$; it does not show that every discrepancy between model and simulation has that origin. The ratio of the decay rate to ϵ in the event-aware solver was 1.93, 1.98 and 1.99 for starts at $0.06T$, $0.02T$ and $0.005T$, independent of ϵT between 0.01 and 0.16 to within 0.01, so it tends to the tangent value 2 as the amplitude falls. The first harmonic predicts 0.116. For this cycle the in-phase state is stable and the antiphase state

is unstable, with $G'(T/2) = +0.77$; Chapter 33 shows that this changes for a cycle with a different charge-to-discharge ratio.

30.3 A Half-Period Delayed Connection

The second mechanism is the symmetric gate with a transport delay τ in both directions, for instance a long passage. By Proposition 30.1 it keeps the equilibria 0 and $T/2$ and changes their stability. Evaluating $G'_\tau(0)$ and $G'_\tau(T/2)$ over τ shows that antiphase is stable and in-phase unstable for $0.34 < \tau/T < 0.80$, with the strongest antiphase stability at $\tau/T = 0.512$, where $G'(T/2) = -1.96$ and $G'(0) = +0.81$. At the edges of the window the stability rate falls to a small fraction of that value, for instance $G'(T/2) = -0.28$ at $\tau/T = 0.34$.

Locking range. With $\tau = 0.512T$ and $\epsilon = 0.08$ the odd function G_τ has extreme values ± 0.2145 around the antiphase branch, so locked antiphase operation exists for $|\delta| \le 0.0172$. Hybrid runs with supply pressure differences of ± 0.3, ± 0.2 and ± 0.1 gave $|\delta| = 0.036$, 0.024 and 0.012; the first two did not lock and the last locked, as predicted.

Offset under detuning. For $|\delta|$ inside the range the locked offset moves from $T/2$ by approximately $\delta/(\epsilon|G'(T/2)|)$, which is 6.4 time units per unit of δ here. Solving $\delta + \epsilon G_\tau(\varphi) = 0$ gave offsets $\varphi/T = 0.396$ and 0.602 for $\delta = \mp 0.012$, and the event-aware simulation gave 0.396 and 0.602, with 0.500 at $\delta = 0$. Simulated offsets are phase offsets in cycles of the locked period, that is, one minus the lag between valve closings divided by the period of the reference chamber; dividing the lag by the nominal period instead biases the offset by the product of the relative frequency difference and the offset, up to 0.006 at these detunings, and that convention gave 0.388 and 0.601 in an earlier comparison.

Sensitivity to period. Because the delay is a time, a change of the period moves τ/T. For a design delay of $0.512T_0$ the window corresponds to periods between about $0.64T_0$ and $1.5T_0$, but the stability rate shrinks toward the edges.

30.4 A Check-Valve Connection

The third mechanism is a one-way gate that passes flow only from chamber 1 to chamber 2 when $p_1 > p_2$, with $q_2 = g[p_1 - p_2]_+$ and $q_1 = -q_2$ by conservation. The rectification breaks the symmetry, so stable offsets other than 0 and $T/2$ become possible. Equation (30.1) holds with $G = Q$,

$$Q(\varphi) = \Big\langle \big(Z(s) + Z(s+\varphi)\big)\,\big[p(s) - p(s+\varphi)\big]_+ \Big\rangle .$$

Q has mean 0.073, first-harmonic amplitude 0.123 and extreme values -0.052 and 0.186, with a first-harmonic fit error of 5 percent of its range. Locked states satisfy $Q(\varphi) = -\delta/\epsilon$ and are stable where $Q' < 0$, which holds for φ/T in $(0.6625, 1)$ and in $(0, 0.157)$. The locking range is $-0.186\epsilon \le \delta \le 0.052\epsilon$, which is $[-0.0149, 0.0042]$ at $\epsilon = 0.08$, and the realizable offsets, taken modulo one period, form the arc from $\varphi/T = 0.6625$ through 0 (equivalently 1) to $\varphi/T = 0.157$, of width 0.495. Antiphase lies in the unstable interval and is not realizable by this valve alone.

Hybrid simulation with $\epsilon = 0.08$ and detuning by supply pressure agreed with the predicted stable offsets: for $\delta = -0.0119$ the prediction was 0.8139 and the simulation 0.8149, for $\delta = -0.0059$ it was 0.9249 and 0.9262, and for $\delta = 0.0024$ it was 0.0519 and 0.0504. Cases outside the range, $\delta = -0.0179$ and 0.0094, did not lock, and $\delta = 0.0047$, just outside the predicted range 0.0042, was slipping slowly. This mechanism therefore realizes a tunable offset in a bounded arc by choosing the detuning, and the tuning range is narrow, about two percent in relative rate.

30.5 Sensitivity to Load

A bladder load on a chamber shifts its frequency, and to first order the relative rate shift is $\delta_{\mathrm{load}} = (g_b/C)\langle Z(p_b - p)\rangle$, with p_b the periodic steady-state bladder response of Chapter 24, which has to be computed at its periodic steady state: the bladder time constant $1/[(1+\nu)g_b]$, with ν the vent ratio, is 133 time units for the smallest load below, and a response started from a fixed pressure and run for fewer cycles overstates the shift. For a bladder with $C_b = C$, vent conductance half the support conductance, and $g_b/C = 0.005$, 0.01, 0.02, the shifts are -0.0008, -0.0016 and -0.0031, which use 5, 9 and 18 percent of the antiphase locking range 0.0172 when only one chamber is loaded. A hybrid simulation of a single loaded chamber, with the gate closed and the bladder started at its periodic mean, gave -0.0008, -0.0016 and -0.0032 for the same three loads. Loads of this size are tolerable, and larger or unequal loads consume the range directly and move the antiphase offset by $\delta_{\mathrm{load}}/(\epsilon|G'(T/2)|)$. These values are first-order in g_b/C and are specific to the cycle used.

30.6 Edges for a Gait Network

The realizable edges are therefore an in-phase edge by the diffusive gate, with local strength g/C; an antiphase edge by the half-period delayed gate, within a bounded detuning range; and a tunable-offset edge by the check valve within a narrow arc of offsets. With edges restricted to in-phase and antiphase, frustration has a concrete form.

Proposition 30.3 (Parity of a triangle)**.** *A triangle of chambers whose three edges are each in-phase or antiphase has consistent offsets if and only if an even number of its edges are antiphase.*

Proof. The offsets around the triangle sum to $k\pi$ with k the number of antiphase edges, and consistency requires the sum to vanish modulo 2π. □

Three groups that are all mutually antiphase are therefore frustrated by construction. The stance, free-leg and arms example has a consistent assignment, for instance the two legs in antiphase and the arms in phase with the stance leg, which has two antiphase edges around the cycle. The general frustrated case is left for later.

Chapter 31

From Chamber Phases to Cord Lengths

In the first embodiment the chambers are the oscillators and the drums are commanded. The coordination result of Part III then concerns the pneumatic cycles, and coordinated tendon motion depends on an interface that this chapter specifies.

31.1 The Mapping

The blackbox senses the phase of the chamber assigned to each channel and commands the drum from it. With phases θ_i in angle units,

$$\ell_c^{\mathrm{cmd}}(t) = \bar{\ell}_c + A_c\, s\big(\hat{\theta}_{i(c)}(t) + \chi_c\big), \tag{31.1}$$

where $i(c)$ is the chamber assigned to channel c, s is a 2π-periodic waveform such as the cosine, A_c is the stroke, χ_c is the intended phase of the cord's motion relative to its chamber, and $\hat{\theta}$ is the phase the blackbox estimates. The estimate comes from the valve events and an interpolation or model between them, as in Chapter 36, and its error is assumed bounded by e_s.

31.2 How the Take-Up Follows

The drum and its actuator are modeled as a first-order lag with time constant τ_c, $\tau_c \dot{q}_c = \ell_c^{\mathrm{cmd}} - q_c$, where τ_c depends on the load. For a sinusoidal command at angular frequency ω the steady-state drum position is

$$q_c = \bar{\ell}_c + \frac{A_c}{\sqrt{1+\omega^2\tau_c^2}}\cos\big(\hat{\theta}_{i(c)} + \chi_c - \zeta_c\big), \qquad \zeta_c = \arctan(\omega\tau_c), \tag{31.2}$$

so the drum lags its command by ζ_c and its amplitude is reduced. The lag is a systematic phase error that can be compensated by choosing χ_c to include a lead $\hat{\zeta}_c = \arctan(\omega\hat{\tau}_c)$ for a nominal time constant $\hat{\tau}_c$. The residual error is $\hat{\zeta}_c - \zeta_c$, which satisfies $|\hat{\zeta}_c - \zeta_c| \le \omega|\hat{\tau}_c - \tau_c|$ because the arctangent has derivative at most one.

Proposition 31.1 (Tendon-layer coordination)**.** *Let the pneumatic phases of group a have pattern-adjusted spread at most γ_a, let channel phases be chosen with the lead compensation above, and let the estimate error satisfy $|\hat{\theta} - \theta| \le e_s$. Then in steady sinusoidal operation the drum phases of the channels of group a, adjusted by their intended offsets, have spread at most $\gamma_a + 2e_s + 2\omega \max_c |\hat{\tau}_c - \tau_c|$.*

Proof. The adjusted phase of the drum of channel c is $\theta_{i(c)} + (\hat{\theta} - \theta) + (\hat{\zeta}_c - \zeta_c)$ plus the common intended pattern. The spread of the first term is at most γ_a by hypothesis, the second contributes at most $2e_s$ to a spread, and the third at most $2 \max_c |\hat{\zeta}_c - \zeta_c| \le 2\omega \max_c |\hat{\tau}_c - \tau_c|$. □

The result is an error budget and has three limits that should be stated. It assumes a constant frequency and a slowly varying phase, so that the steady-state lag formula applies, and a first-order drum, whereas a drum with saturation of speed or torque, or higher-order dynamics, needs its own analysis. It ends at the drum coordinate: how a drum position becomes a joint motion depends on routing, slack, elasticity and load, and no statement about joint angles follows without a model of that map. And the load enters through τ_c, so a load that changes the time constant changes the lag and the amplitude, which the compensation handles only to the accuracy of $\hat{\tau}_c$. The drum motion also feeds back on the pneumatic layer only through the body path of Chapter 24, which is already modeled as a bounded perturbation.

Chapter 32

A Worked Three-Chamber Gait Network

This chapter combines the connections of Chapter 30 in the consistent assignment identified earlier, loads the chambers, connects them to drum commands through the interface of Chapter 31, and runs a release and recapture. Everything here is a calculation within the model and a simulation of the model: the cycle is the one of Chapter 30, the gates and delays are ideal, the loads are linear bladders, the drums are first-order, and no joint mechanics are included. No claim reaches the physical body.

32.1 The Network

Three chambers represent the stance group S, the free leg F and the arms A. The stance and arm chambers are joined by a diffusive gate, an in-phase edge, with $\epsilon_{SA} = 0.06$. The free-leg chamber is joined to each of them by a delayed gate, an antiphase edge, with delay $\tau = 0.512T$ and $\epsilon_{SF} = 0.06$, $\epsilon_{AF} = 0.05$. The target is stance and arms in phase, free leg in antiphase, which is consistent by the parity rule. Each chamber carries a bladder with $C_b = C$ and vent conductance half the support conductance, with $g_b/C = 0.02$, 0.01, 0.005 for S, F, A, giving first-order load shifts of the natural rate $\delta_S = -0.0031$, $\delta_F = -0.0016$, $\delta_A = -0.0008$ from Chapter 30. A single loaded chamber in the hybrid simulation, with the gates closed, shifted by -0.0032, -0.0016 and -0.0008, which agrees.

Two kinds of number appear below and are kept apart. Averaged-model values are solutions of the first-order equations of this chapter. Hybrid values come from simulations in which the valves switch at the thresholds and the flows act continuously; they are reported from an event-aware fourth-order solver that locates each threshold crossing and restarts the integration at the switch, with the bladders started at their periodic mean pressure. A forward-Euler solver at the step 2.5×10^{-5} agrees with it to the digits quoted unless stated, and the dependence of the hybrid values on the solver, the step and those choices is documented in Chapter 37. Hybrid offsets are phase offsets in cycles of the locked period, equal to one minus the lag between a valve closing of the reference chamber S and the next closing of the chamber in question, divided by the period of S. They agree to 10^{-4} with the instantaneous phase offsets θ_i/T_i read from the pressures at the closing of S. The averaged model computes cycle-averaged offsets, which differ from the instantaneous ones by the oscillatory

displacement $W_i - W_S$ of Chapter 34, up to 0.0004 in the cases tested; the comparisons below are between the averaged offsets and the instantaneous hybrid offsets, and the agreement quoted includes that difference unless a correction is stated.

32.2 Equations and the Locked Solution

With the interaction function derived there, $\bar{\Gamma}(x) = \langle Z(s)[p(s+x) - p(s)]\rangle$, the averaged phases obey

$$\dot{\theta}_i = 1 + \delta_i + \sum_j \epsilon_{ij}\, \bar{\Gamma}\big(\theta_j - \theta_i - \Omega\tau_{ij}\big), \tag{32.1}$$

in time units, where Ω is the collective rate and $\Omega\tau_{ij}$ is the phase advanced during the delay. The delay is a time, so its phase equivalent depends on the collective rate, which is why Ω appears and must be solved for with the offsets. The shift of the argument by $\Omega\tau_{ij}$ treats the phase as slowly varying during the delay, a quasi-static approximation that the hybrid simulation tests. The function $\bar{\Gamma}$ is tabulated from the cycle and no sinusoidal approximation is used, so no sinusoidal fitting error enters a remainder budget, although the first-order averaging error remains (the phase equation itself is exact, Theorem 34.1) and is examined in Chapter 34; the sinusoidal three-group theorem is not invoked, and stability is established directly for this network by linearization.

A common locked solution has $\dot{\theta}_i = \Omega$ for all i, which gives three equations for Ω, $u_F = \theta_F - \theta_S$ and $u_A = \theta_A - \theta_S$. One consequence is visible before solving. The free-leg chamber has two delayed gates and each of the others has one, so the constant part of $\bar{\Gamma}$ at the antiphase argument, $\bar{\Gamma}(T/2 - \tau) = -0.0107$, shifts the free-leg rate by twice as much as the others, an effective detuning of $\epsilon\,\bar{\Gamma}$ of order 6×10^{-4} that the locked offsets must absorb. It is small here.

Case	Ω	u_F/T	u_A/T	Relaxation rates
No loads, averaged model	0.9992	0.4965	0.0009	0.157, 0.177
Loaded, averaged model	0.9975	0.5065	0.0160	0.171, 0.156
Loaded, hybrid simulation	0.9975	0.5064	0.0158	0.17, 0.15 (fitted)

The locked solution exists and is linearly stable in the averaged model, with relaxation rates near 0.16 per unit time, and the hybrid simulation relaxes at rates of the same size. The fitted hybrid rates, obtained by displacing F or A by $0.04T$ from the locked state, vary by about ± 0.01 with the step and with the fitting window, so agreement is claimed only to that level. The averaged and hybrid locked states differ by $0.0001T$ in u_F, $0.0002T$ in u_A and by less than 0.0001 in Ω, which is within the step-size variation of the hybrid values, $0.0005T$ for steps of 10^{-4} and below, so no discrepancy between the first-order averaged model and the hybrid simulation is resolved for this network at this coupling. The loads move the offsets by 0.010 and 0.015 of a period and the collective rate by 0.2 percent.

32.3 Margin After Loading

The pairwise locking range of a delayed edge is $\epsilon \max|G_\tau| = 0.2145\epsilon$, which is ± 0.0129 for $\epsilon = 0.06$ and ± 0.0107 for 0.05, and it does not establish that the three edges admit a common

solution. The network was therefore solved with extra detuning added to one chamber at a time, starting from the loaded state, and the solution was continued until it ceased to exist or lost stability.

Extra detuning of chamber	Lower margin	Upper margin
S	-0.0148	$+0.0233$
F	-0.0156	$+0.0199$
A	-0.0169	$+0.0178$

When all three chambers are detuned together, with every combination of signs, the largest common magnitude for which all eight corners remain stable was 0.0075; this is a numerical check on the corners of a box. It remains a numerical check unless an argument, for instance a monotonicity or convexity property of the locked solution in the detunings, shows that the corners cover the whole box, and none is given. The limits of locking of the hybrid simulation were found for extra detuning of F by supply pressure from runs of 3000 time units, bracketing the limit between the largest detuning that stayed locked and the smallest that slipped, and independently by ramping the detuning slowly and extrapolating the loss point to zero ramp rate as $r^{2/3}$. Both methods gave -0.0151 and $+0.0198$, against -0.0156 and $+0.0199$ from the averaged model. Fixed-time bisection with runs of 400 to 500 time units gives different limits for a delayed pair because a slip near the fold takes longer than the run to appear, so that method is not used (Chapter 37).

Extra δ_F	Predicted u_F/T	Simulated u_F/T	Simulated state
$+0.010$	0.5824	0.5812	locked
$+0.016$	0.6503	0.6487	locked
$+0.024$	none	0.77, spread 0.025	not locked
-0.010	0.4267	0.4269	locked
-0.020	none	0.22, spread 0.043	not locked

The simulated offsets differ from the averaged prediction by at most $0.0016T$ among the locked cases. The spread is the circular standard deviation of the offset over the last twenty cycles, in units of T.

The delayed edge's effective phase depends on the collective period through $\Omega\tau/T$. Across the locked cases above, Ω varies between 0.994 and 1.002 and $\Omega\tau/T$ stays within 0.509 and 0.513, far from the window 0.34 to 0.80. A change of the period by several tens of percent would be needed to approach the window edge, where the stability rate collapses.

32.4 Supply Sharing

When the chambers draw from a shared line with resistance r_{line}, so that the supply seen by each charging chamber drops by r_{line} times the total charging flow, the network stayed locked but its offsets moved.

r_{line}	Ω	u_F/T	u_A/T	Shift of u_F, u_A
0	0.9975	0.5064	0.0158	–
0.005	0.9907	0.5193	0.0149	$+0.0129$, -0.0009
0.01	0.9839	0.5360	0.0147	$+0.0296$, -0.0011
0.02	0.9693	0.6000	0.0148	$+0.0936$, -0.0010

The shared line acts as a state-dependent coupling between the chambers, since which chambers are charging at once depends on their relative phases. A bounded perturbation can represent it provided a bound valid uniformly over the states is established, and that has not been done here. What such a representation discards is the structure, namely which chambers charge together, so a guarantee obtained from it would be conservative, and it could not predict the direction of the offset shifts. The shared line moves the free leg relative to the other two by 0.013 to 0.094 of a period, more than the loads do at the larger resistances, and leaves the offset of the retained pair S, A almost unchanged, to within $0.0011T$.

32.5 Release and Recapture

The free leg was released by closing both of its gates to a residual conductance of 2 percent and detuning its supply so that $\delta_F = +0.03$, beyond any locking range the residual coupling allows. The release lasted from time 60 until 160, 168 or 176, after which the gates were restored and the detuning removed. The retained pair S, A stayed coupled throughout.

Release ends at	160	168	176
u_F/T before release	0.5064	0.5064	0.5064
u_F/T at recapture	0.8375	0.1168	0.3914
u_F/T after recapture	0.5064	0.5064	0.5064
Settling time to within $0.02T$	23	30	10
Net unwrapped change in u_F/T	-0.329	$+0.406$	$+0.106$

The free leg returned to the same antiphase offset to within $0.0001T$ from three different phases at recapture. Three cases do not show that this holds from every phase, so the release end was scanned at intervals of half a time unit over thirty time units, which carries the recapture phase around the whole cycle, and each of the 61 runs was continued for four hundred time units. The free leg recaptured to the intended state from every tested initial condition: all 61 ended at $u_F/T = 0.5064$ and $u_A/T = 0.0158$. The settling time ranged from 0 to 51 time units with a median of 20, and exceeded 40 for four runs, all with recapture phases within about $0.07T$ of the in-phase state, which is unstable for the delayed edge. Three of the runs had a net change of more than half a period, 0.534 at most, and none a full period. A start exactly on the unstable state would not be recaptured. The scan was run with the event-aware solver and with Euler at 2.5×10^{-5}; the final states and classes were identical, and the settling times agreed to 0.06 time units.

The retained pair did shift. Relative to its settled value before release, $u_A/T = 0.0157$, it moved to 0.0228 during release, a shift of $+0.0071$ of a period, and returned to 0.0158 after recapture. The pair is not isolated from the third chamber: removing the free leg changes the frequency difference seen by S and A, which are coupled to it with different strengths and loads. A shift of the retained pair of under one percent of a period during release is therefore an unintended effect of the release, small here, and it should be bounded in any claim that the retained pair keeps its coordination.

32.6 Commanded Cord Lengths

Each chamber commands one drum through the mapping of Chapter 31, with intended tendon offsets equal to the chamber offsets, so that the free-leg drum is in antiphase with the stance drum. The nominal drum time constant is $\hat{\tau} = 0.05$ and the actual values are 0.06, 0.045 and 0.04 for S, F, A, which at $\omega = 6.92$ give lag errors of -0.060, $+0.031$ and $+0.063$ radians after compensation. The blackbox estimates phase by linear interpolation from the last two valve closings, and its error relative to the pressure-based phase was at most 0.014 radians.

Window	Drum error F–S	Bound	Drum error A–S	Bound
Locked, before release	0.132	0.209	0.223	0.268
After recapture, release ended 160	0.132	0.209	0.223	0.269
After recapture, release ended 168	0.132	0.209	0.223	0.269
After recapture, release ended 176	0.132	0.209	0.223	0.269

Errors are drum phase differences from the intended offset in radians, and the bound is the chamber offset error plus twice the estimator error plus twice ω times the largest time-constant mismatch, the pairwise form of Proposition 31.1. Every measured error lies inside its bound, with F–S at about 63 percent of it and A–S at about 83 percent. The simulated drum amplitudes were 0.926, 0.957 and 0.966, against 0.924, 0.955 and 0.964 predicted from the lag. The drum phases of the retained pair and of the recaptured free leg are therefore within about 0.13 and 0.22 radians of their targets, with most of the error coming from the time-constant mismatch and the residual chamber offsets, which are 0.042 and 0.102 radians of the pressure-based phase for F–S and A–S.

The boundary is stated where it lies. The results end at the drum coordinate. The map from a drum position to a joint motion depends on routing, slack, elasticity and load, and no statement about joint angles, balance or support is made. The first-order theory used ϵ at most 0.06, the cycle is one parameter set, and the delay is an ideal transport.

Chapter 33

A Second Cycle

The worked network of Chapter 32 used one pneumatic cycle. This chapter repeats the analysis for a cycle with a substantially different charge-to-discharge ratio, to separate two questions: how far the behavior tolerates a change of the cycle, and whether it can be reproduced after calibration. All results are calculations within the model and simulations of the model, at first order in the coupling, and the first-order averaging error remains (Chapter 34).

33.1 The Two Cycles

The first cycle, A, has $p_s = 10$, $p_h = 6$, $p_l = 3$, $p_o = 0$, $\tau_c = 1$, $\tau_o = 0.5$, a period $T = 0.906$ and a charge-to-discharge time ratio of 1.6. The second cycle, B, has $p_s = 9$, $p_h = 6.5$, $p_l = 2.5$, $p_o = 0$, $\tau_c = 1.6$, $\tau_o = 0.25$, a period $T = 1.768$ and a ratio of 6.4: a slow charge and a fast discharge. The waveform of B is a slow rise with an abrupt fall, and the phase response is correspondingly concentrated.

33.2 Interaction Functions

The normalization $Z\dot{p} = 1$ still holds, so $G'(0) = -2$ for the diffusive gate in both cycles to the accuracy of the numerics. Everything else in the interaction function changes.

Diffusive gate	Cycle A	Cycle B
$G'(0)$	-2.00	-2.00
$G'(T/2)$	$+0.76$	-0.18
Stable states	in-phase only	in-phase, and antiphase weakly
Sine coefficients, harmonics 1 to 3	-0.210, -0.014, -0.002	-0.112, -0.070, -0.036
Peak error of first-harmonic fit	8.5 percent	64 percent
Slope at 0 from first harmonic	1.45	0.40
Peak of G against first-harmonic amplitude	0.208 against 0.210	0.175 against 0.112

For A the sinusoid was accurate for the locking range and poor for the local slope. For B it is poor for both: the second harmonic is more than half the first, the slope at zero is underestimated by a factor of five, and the locking range is underestimated by more than a third. The sinusoidal approximation therefore cannot be carried from one cycle to the other,

which is the case for the direct treatment of the derived function in the worked network. For B the antiphase state of the diffusive gate is also weakly stable, so a single diffusive gate no longer selects in-phase.

For the delayed gate the window in which antiphase is stable and in-phase unstable moves.

Delayed gate	Cycle A	Cycle B
Window in τ/T	0.34 to 0.80	0.58 to 0.92
At $\tau = T/2$: $G'(0)$, $G'(T/2)$	$+0.76$, -2.00	-0.18, -2.00
Delay of $0.512T_A = 0.464$, as a fraction of T	0.512	0.2625
$G'(0)$, $G'(T/2)$ at that delay	$+0.81$, -1.96	-0.58, $+0.60$
Further zeros of G appear for τ/T near	0.31 to 0.33, 0.81 to 0.83	0.42 to 0.44, 0.92 to 0.94

Keeping the original physical delay, 0.464 time units, places it at 0.2625 of the new period, outside the window and in the regime where antiphase is unstable and in-phase is stable. Resetting the delay to half the new period puts antiphase at its strongest stability but leaves in-phase stable as well, for the pair. The one-way gate was not simulated for B; its computed interaction function ranges from -0.003 to 0.739, which is strongly one-sided.

33.3 The Network Under Four Settings

The network of Chapter 32 was recomputed for B with the couplings scaled by the ratio of periods, so that ϵT is as in cycle A, and the bladder conductances scaled the same way, which gives first-order load shifts -0.0071, -0.0036, -0.0018 for S, F, A. Four settings were compared by solving for the stable states of the averaged model, by the box margin of the averaged model, and by a hybrid simulation started near the intended antiphase state and run for 3000 time units with the event-aware solver, long enough for states that attract at rates near 0.015 to settle.

Setting	Hybrid u_F/T, u_A/T	Averaged stable states	Box margin
Original physical delay 0.464	0.221, 0.347	(0.220, 0.355), rate 0.015	–
Delay reset to $T/2$	0.535, 0.048	(0.531, 0.044)	0.0009
Delay $0.72T$, from the pair window	0.663, 0.276	(0.654, 0.271), rate 0.018	–
Retuned: delay $T/2$, couplings $\times 3$	0.509, 0.013	(0.508, 0.012)	0.0071

Cycle A had a box margin of 0.0075. The first row fails: the network ends far from the intended antiphase state, with the free-leg chamber at $0.22T$ from the stance chamber and the arms at $0.35T$, and the averaged model has a weakly stable state close to it. These states attract at a rate of 0.015 per unit time, so a run of 260 time units, which is long enough for the other settings, shows only a transient: at that time the same setting read $u_F/T = 0.025$ and $u_A/T = 0.179$, and the third setting read 0.554 and 0.178, against 0.221, 0.347 and 0.663, 0.276 after 3000 time units. The averaged and hybrid states agree to $0.001T$ and $0.008T$ in the first row and $0.009T$ and $0.005T$ in the third. The second row works but with

a box margin about eight times smaller than in A. The third shows that the pair criterion, in-phase unstable and antiphase stable, does not guide the network: the delayed gates then differ from the antiphase argument by $T/2-\tau$, the constant part of the interaction there is $\bar{\Gamma}(T/2-\tau)=-0.20$, and the free-leg chamber, which has two delayed gates, is shifted twice as much as each of the others, an effective relative detuning of about 0.006 that is comparable to the pairwise locking range 0.0063. At $\tau=T/2$ the constant term is $\bar{\Gamma}(0)=0$, so the imbalance vanishes and the unloaded network has the exact antiphase solution $u_F=T/2$, $u_A=0$, $\Omega=1$. This is the reason the half-period delay works at the network level, and it is an observation about this cycle and not a rule.

33.4 The Retuned Network

For the retuned network the unloaded averaged solution is exactly antiphase and the loaded solution has $\Omega=0.9961$, $u_F=0.5084T$, $u_A=0.0119T$ with a doubled relaxation rate 0.234 per unit time; the hybrid simulation gave 0.9960, $0.5094T$, $0.0131T$ and a fitted rate of 0.256 ± 0.004 after displacing F by $0.04T$, so the offsets differ by $0.0010T$ and $0.0012T$. The load shifts use 44, 22 and 11 percent of the pairwise locking range of the delayed edge, against 24, 12 and 6 percent in cycle A, so the period-scaled loads are a larger fraction of the range in B. The margins for extra detuning of one chamber, from the loaded state, were $[-0.0135,+0.0349]$ for S, $[-0.0174,+0.0258]$ for F and $[-0.0197,+0.0242]$ for A in the averaged model, with a corner-check box margin of 0.0071; as before the corner check is numerical.

Here the averaged model is noticeably less accurate than in A. The limits of locking of the hybrid simulation for extra detuning of F, from runs of 3000 time units, are -0.0140 and $+0.0207$, narrower than the averaged limits by 0.0034 and 0.0051, that is by 20 percent of the averaged limit at each end (the prediction P1 of Chapter 34, against the hybrid with dynamic bladders). Chapter 34 diagnoses this discrepancy. Three corrections were tested and are insufficient to account for it: freezing or removing the bladder dynamics changes the limits by at most 0.0002, recomputing the interaction function for the detuned chambers removes only a few percent, and the second-order averaging term as computed does not reproduce the limits. When the frequency spread of the chambers is scaled with the coupling the discrepancy falls roughly in proportion to it, as a first-order averaging error would, and it does not fall when the spread is held fixed while the coupling is reduced, which leaves the regime of the theory. That trend is consistent with an averaging error but does not by itself identify the cause.

Extra δ_F	Predicted u_F/T	Simulated u_F/T	Simulated state
$+0.010$	0.5344	0.5374	locked
$+0.020$	0.5715	0.5894	locked, near the limit 0.0207
$+0.030$	none	0.63, spread 0.064	not locked
-0.010	0.4817	0.4800	locked
-0.020	none	0.67, spread 0.21	not locked

The case $+0.020$ lies within four percent of the hybrid limit, where the relaxation is slow. Its simulated offset differs from the prediction by $0.018T$, the largest of the locked cases.

33.5 Release, Recapture and Drum Tracking

The release closed both free-leg gates to a residual conductance of 2 percent and detuned the free-leg supply to $\delta_F = +0.04$, beyond the pairwise range scaled as before, from time 80 until 160, 175 or 190. Each run was continued for four hundred time units after the gates were restored, because the settling is slow from some phases.

Release ends at	160	175	190
u_F/T before release (window 60 to 80)	0.5094	0.5094	0.5094
u_F/T at recapture	0.3485	0.7139	0.0409
u_F/T four hundred time units later	0.5074	0.5074	0.5074
Settling time to within $0.02T$	139	20	98
Net unwrapped change in u_F/T	$+0.154$	-0.206	$+0.443$

The free leg recaptured to antiphase from all three phases, ending at $0.5094T$, the pre-release value, with no full-cycle slip. The settling times are very different from case to case, 20 to 139 time units against 10 to 30 in cycle A, and the case ending at 160 sat for a long stretch near an alternative state with the free leg in phase with the stance chamber and the arms in antiphase to both. To see how general this is, the release end was scanned at intervals of half a time unit over forty-five and a half time units, which carries the recapture phase around the whole cycle, giving 92 runs, each continued for four hundred time units. The free leg recaptured to the intended state from every tested initial condition: all 92 ended at $u_F/T = 0.5094$ and $u_A/T = 0.0131$. The settling time ranged from 0 to 142 time units with a median of 39. It was under 20 for 33 runs, all with recapture phases between $0.38T$ and $0.71T$, and above 90 for 37 runs, all with recapture phases between $0.95T$ and $0.35T$ through the in-phase state. Seven of the 92 had a net change above half a period, 0.556 at most, and none a full period. Recapture is therefore complete in the sense that every tested start returned to the intended state, but it is fast only for starts within roughly a third of a period of antiphase. The reason is visible in the pair interaction of the table above: at $\tau = T/2$ in this cycle the in-phase state of a pair has $G'(0) = -0.18$, so the neighborhood of in-phase is nearly flat, and a free leg released into it leaves slowly. This explanation was not tested separately. The scan was run with the event-aware solver and with Euler at 3.75×10^{-5}. The final states agreed in all runs. The settling times agreed to 0.03 time units at the median, but one run, with recapture phase $0.94T$, near the boundary where the route around the cycle reverses, went the opposite way round in the two solvers (net change $+0.556$ against -0.444, settling times 113 and 67), so the location of that boundary is solver-dependent at the level of $0.001T$ in the recapture phase.

The retained pair moved from $u_A/T = 0.0131$ before release to 0.0200 during release, an unintended shift of $+0.007$ of a period, comparable to cycle A, and returned to 0.0131 afterward.

The drum commands used the same physical drum time constants as before, 0.06, 0.045 and 0.04, with the nominal 0.05 and lead compensation recomputed for the new frequency $\omega = 3.54$. The lag errors after compensation are -0.034, $+0.017$ and $+0.035$ radians, smaller than in A because the frequency is lower, and the estimator error was at most 0.014 radians.

Window	Drum error F–S	Bound	Drum error A–S	Bound
Locked, before release	0.110	0.154	0.151	0.176
After recapture, release ended 160	0.110	0.154	0.151	0.176
After recapture, release ended 175	0.110	0.153	0.151	0.176
After recapture, release ended 190	0.110	0.153	0.151	0.176

Every measured error lies inside its bound, with F–S at about 71 percent of it and A–S at about 86 percent. The simulated drum amplitudes were 0.980, 0.989 and 0.991, against 0.978, 0.988 and 0.990 predicted from the lag.

33.6 What Is Retained and What Is Retuned

Retained: the topology and edge types, the consistent assignment, the structure of the interface from chamber phase to drum command, the normalization $G'(0) = -2$ of the diffusive gate, the method of release by gate closure and supply detuning, the physical drum time constants, and the qualitative sequence of lock, release, recapture and tracking. Retuned: the delay, which has to be reset to half the new period and not retained as a physical value; the gate conductances, which had to be raised so that ϵT rose by a factor of three to restore the margin; the release detuning, rescaled to the new locking range; and the lead compensation, recomputed from the new frequency. The pairwise window for the delay is not a reliable guide for the network, and the half-period delay is justified by the vanishing of the constant term in the delayed interaction at the antiphase argument.

The two readings of the question separate cleanly. Tolerance to the change of cycle is poor for the physical delay, which fails outright, and moderate for the structure, which survives. Reproducibility after calibration holds with three adjustments, with two reservations. The coupling has a first-order accuracy cost since ϵT is three times larger, and the averaged model then overestimates the limits of locking of the loaded network by about twenty percent of its own limits (the prediction P1; fifteen to seventeen percent if the interaction is recomputed for the detuned chambers), a discrepancy whose cause Chapter 34 does not establish, since the corrections tested there are insufficient to account for it. And recapture, although it reached the intended state from every tested initial condition, is much slower from phases near in-phase than it was in cycle A. One second cycle does not establish how the required recalibration depends on the cycle ratio, and no general rule is claimed.

Chapter 34

Error of the Averaged Chamber Model

The ledger has listed the distance between the oscillator model and the chambers as the largest unbounded gap. For the chamber model used in this Part that gap has two components, and one of them vanishes. This chapter proves that the phase equation of a single-pressure chamber is exact, bounds the other component, the error of averaging, with explicit constants (a finite-horizon bound, and a conditional persistence statement whose hypotheses fail at the couplings used) that turn out to be uninformative, diagnoses the largest discrepancy found in simulation, which is in cycle B, and states what a useful bound will have to cover. The scope is single-pressure chambers joined by linear gates with constant delays. Chambers with bladder loads have a second state variable; Section 34.3 derives their exact equations and measures the effect of the extra variable, and the averaging theorem is stated for the single-pressure chambers, and its extension to the loaded form is open. The finite-horizon bound is first order in ϵ' at a fixed time; the persistence statement is conditional and its hypotheses are not met at the couplings used.

34.1 The Model and the Exact Phase Equation

Chamber i has a pressure p_i and a mode $s_i \in \{0, 1\}$. In mode 0, charging, $\dot{p}_i = (\sigma_i - p_i)/\tau_c + q_i$, with supply $\sigma_i > p_h$. In mode 1, discharging, $\dot{p}_i = (p_o - p_i)/\tau_o + q_i$. The mode changes from 0 to 1 when $p_i = p_h$ and from 1 to 0 when $p_i = p_l$, with $p_o < p_l < p_h < \sigma_i$. The gate inflow is $q_i(t) = \sum_j \epsilon_{ij}\big(p_j(t-\tau_{ij}) - p_i(t)\big)$ with $\epsilon_{ij} \geq 0$ and $\tau_{ij} \geq 0$ constant. The supplies σ_i may differ between chambers. Write $a = p_h - p_l$ and $t_{c,i} = \tau_c \ln \frac{\sigma_i - p_l}{\sigma_i - p_h}$, $t_d = \tau_o \ln \frac{p_h - p_o}{p_l - p_o}$ and $T_i = t_{c,i} + t_d$.

Assumption W. Initially every $p_i \in [p_l, p_h]$, and for every i

$$a \sum_j \epsilon_{ij} < \min\Big(\frac{\sigma_i - p_h}{\tau_c},\ \frac{p_l - p_o}{\tau_o}\Big). \tag{34.1}$$

For cycle A this requires $\sum_j \epsilon_{ij} < 1.33$ and for cycle B $\sum_j \epsilon_{ij} < 0.39$ per unit time, which holds for every coupling used in this monograph.

Lemma 34.1 (Monotone cycles)**.** *Under Assumption W every pressure stays in* $[p_l, p_h]$*, increases strictly in mode* 0 *and decreases strictly in mode* 1*, and the switching is transversal, so each solution has finitely many switches in bounded time.*

Proof. Suppose all pressures lie in $[p_l, p_h]$ on an interval. In mode 0 with $p_i \in [p_l, p_h)$, $q_i \geq -\sum_j \epsilon_{ij}(p_i - p_l) \geq -a\sum_j \epsilon_{ij}$ and $(\sigma_i - p_i)/\tau_c \geq (\sigma_i - p_h)/\tau_c$, so $\dot{p}_i > 0$ by Assumption W. In mode 1 with $p_i \in (p_l, p_h]$, $q_i \leq a\sum_j \epsilon_{ij}$ and $(p_o - p_i)/\tau_o \leq -(p_l - p_o)/\tau_o$, so $\dot{p}_i < 0$. A chamber in mode 0 therefore rises from p_l to p_h and cannot leave the interval before switching, and likewise in mode 1; since $\dot{p}_i$ is bounded away from zero, each mode lasts a bounded time. The continuation argument closes because the first exit from $[p_l, p_h]$ would have to occur with the wrong sign of $\dot{p}_i$. □

Theorem 34.1 (Exact phase equation)**.** *Under Assumption W define, in mode* 0*,* $\theta_i = -\tau_c \ln \frac{\sigma_i - p_i}{\sigma_i - p_l} \in [0, t_{c,i}]$ *and, in mode* 1*,* $\theta_i = t_{c,i} - \tau_o \ln \frac{p_i - p_o}{p_h - p_o} \in [t_{c,i}, T_i]$*, and let* $\Theta_i = \theta_i + n_i T_i$ *with* n_i *the number of completed cycles. Then* θ_i *is continuous across switches,* $p_i = P_i(\theta_i)$ *with the nominal waveform* P_i *of the cycle of chamber* i*, and*

$$\dot{\Theta}_i = 1 + Z_i(\theta_i)\, q_i(t), \qquad Z_i(\theta) = \begin{cases} \dfrac{\tau_c}{\sigma_i - p_l} e^{\theta/\tau_c}, & 0 \leq \theta < t_{c,i}, \\ -\dfrac{\tau_o}{p_h - p_o} e^{(\theta - t_{c,i})/\tau_o}, & t_{c,i} \leq \theta < T_i, \end{cases} \tag{34.2}$$

exactly, with $q_i = \sum_j \epsilon_{ij}(P_j(\theta_j(t - \tau_{ij})) - P_i(\theta_i(t)))$*. The functions* P_i *and* Z_i *are the periodic waveform and phase response of Chapter 29, so the phase model of that chapter is not an approximation of the chambers.*

Proof. In mode 0, $\partial\theta_i/\partial p_i = \tau_c/(\sigma_i - p_i)$ and the uncoupled vector field is $(\sigma_i - p_i)/\tau_c$, whose product is 1; the chain rule gives $\dot{\theta}_i = 1 + \frac{\tau_c}{\sigma_i - p_i} q_i$, and $\sigma_i - p_i = (\sigma_i - p_l)e^{-\theta_i/\tau_c}$ gives the stated Z_i. In mode 1, $\partial\theta_i/\partial p_i = -\tau_o/(p_i - p_o)$, the uncoupled field is $-(p_i - p_o)/\tau_o$, the product is 1, and $p_i - p_o = (p_h - p_o)e^{-(\theta_i - t_{c,i})/\tau_o}$ gives Z_i. At $p_i = p_h$ both expressions for θ_i equal $t_{c,i}$; at $p_i = p_l$ in mode 1 the second expression equals $t_{c,i} + t_d = T_i$, which is 0 in the next cycle. Lemma 34.1 guarantees that θ_i stays in the stated ranges, so no extension of the formulas is used. □

Equation (34.2) is the unaveraged phase equation, and it is this equation that is exact. The averaged interaction equation of Chapter 29, obtained by replacing its right-hand side by an average over the cycle, is an approximation to it and is not exact. The equation is exact because the chamber has one pressure variable: the phase along the uncoupled flow is a coordinate on the whole state space of each mode, so there is no transverse direction and no remainder of phase reduction. It fails for a loaded chamber, whose bladder adds a state variable and with it a transverse direction. A finite-difference check on a coupled pair gave $|\dot{\theta}_i - (1 + Z_i q_i)|$ with a median of 3×10^{-12} and a maximum of 7×10^{-5}, the maximum arising at the kinks of the sampled waveform. The error of the first-order model of Chapter 29 is therefore the error of averaging (34.2), and it is the object of the rest of this chapter.

34.2 An Averaging Bound for the Exact Equation

Let $f_0 = 1/T_0$ be a base frequency, $x_i = \Theta_i/T_i$ the phase in cycles, $\psi_i = x_i - f_0 t$ and $\tau = f_0 t$. Write $\epsilon_{ij} = \epsilon\kappa_{ij}$ with $\epsilon = \max \epsilon_{ij}$, $\epsilon' = \epsilon T_0$, $K = \max_i \sum_j \kappa_{ij}$, $f_i = 1/T_i$, $\tilde{\delta}_i = (f_i/f_0 - 1)/\epsilon'$,

$\zeta_i(x) = f_i Z_i(xT_i)$, $\pi_i(x) = P_i(xT_i)$ and $d_{ij} = f_0\tau_{ij}$. Then Theorem 34.1 reads

$$\frac{d\psi_i}{d\tau} = \epsilon' F_i(\tau, \psi_\tau), \quad F_i = \tilde{\delta}_i + \sum_j \kappa_{ij}\,\zeta_i(\tau+\psi_i)\big[\pi_j(\tau - d_{ij} + \psi_j(\tau - d_{ij})) - \pi_i(\tau+\psi_i)\big], \tag{34.3}$$

with ζ_i and π_i one-periodic. The functions ζ_i have two jump discontinuities per period, at the two switching phases, and π_i is continuous with two kinks. The averaged field is

$$\bar{F}_i(y) = \tilde{\delta}_i + \sum_j \kappa_{ij}\Gamma_{ij}(y_j - y_i - d_{ij}), \qquad \Gamma_{ij}(w) = \int_0^1 \zeta_i(s)\big[\pi_j(s+w) - \pi_i(s)\big]ds, \tag{34.4}$$

which is the first-order model of Chapter 29 with chamber-specific waveforms, and with the delay phase $f_0\tau_{ij}$ in place of the collective-rate phase $\Omega\tau_{ij}$, a difference of order ϵ'.

Constants. Define, from the cycles: $M = \max_i |\tilde{\delta}_i| + K\zeta_{\max}a$ with $\zeta_{\max} = \max|\zeta_i|$; $\bar{d} = \max d_{ij}$; $\gamma = \max_{i,j,w} |\Gamma'_{ij}(w)|$ and $\Lambda = 2K\gamma$; $q = \max_{i,j,x} \int_0^1 |\zeta_i(s)||\pi'_j(s+x)|ds$ (at least 1, since $|\zeta_i\pi'_i| = f_i$ integrates to 1); $\bar{z} = \max_i \int_0^1 |\zeta'_i|ds$ over the continuous parts; $J = a\max_i$ max jump of ζ_i; $L_1 = K(\bar{z}a + 2q)$; and

$$C = M(1+\bar{d})\big(L_1 + \Lambda + 2KJ\big). \tag{34.5}$$

Theorem 34.2 (Finite-horizon averaging bound)**.** *Under Assumption W, if ψ solves the exact equation for $\tau \geq 0$, with any history on $[-\bar{d}, 0]$ having slope at most $\epsilon' M$, and y solves $dy/d\tau = \epsilon'\bar{F}(y)$ with $y(0) = \psi(0)$, then for all $\tau \geq 0$*

$$\|\psi(\tau) - y(\tau)\|_\infty \leq \big[\epsilon'^2(\tau+1)\,C + \epsilon' M\big]e^{\epsilon'\Lambda\tau}. \tag{34.6}$$

In particular for $\tau \leq L/\epsilon'$ the error is at most $\epsilon'\big[M + (L+\epsilon')C\big]e^{\Lambda L}$. At an integer time N, the end of a complete window, the sharper stroboscopic estimate

$$\|\psi(N) - y(N)\|_\infty \leq \epsilon'^2 NC + \big[\epsilon'^2(N+1)C + \epsilon' M\big]\big(e^{\epsilon'\Lambda N} - 1\big) \tag{34.7}$$

holds. It is of order ϵ'^2 only at a fixed integer time N. Over an averaging horizon $N = L/\epsilon'$ the factor N makes the first term of order ϵ', and the exponential is $e^{\Lambda L}$, which need not be close to one but does not grow as ϵ' decreases, so on the timescale of the averaged flow the estimate is first order, like the general one, and the integer-time restriction gains nothing there.

Proof. Since $|F_i| \leq M$, $|\psi(s) - \psi(s')| \leq \epsilon' M|s - s'|$. Let $N = \lfloor\tau\rfloor$ and cover $[0,\tau]$ by the complete windows $[k, k+1]$, $0 \leq k < N$, and the unfinished window $[N, \tau]$. Fix a window, let $\hat{\psi}$ be the constant function with value $\psi(k)$, and put $D = \epsilon' M(1+\bar{d})$, which bounds $|\psi_j(s - d_{ij}) - \psi_j(k)|$ and $|\psi_i(s) - \psi_i(k)|$ for s in the window. Write $F(s,\psi_s) - \bar{F}(\psi(s)) = [F(s,\psi_s) - F(s,\hat{\psi})] + [F(s,\hat{\psi}) - \bar{F}(\psi(k))] + [\bar{F}(\psi(k)) - \bar{F}(\psi(s))]$. The middle term integrates to zero over a complete window, because $F(\cdot,\hat{\psi})$ is one-periodic with mean $\bar{F}(\psi(k))$. It does not integrate to zero over the unfinished window, and that case is treated separately below. The last term is at most ΛD in sup norm. For the first, split the window into the set where $s + \psi_i(s)$ and $s + \psi_i(k)$ lie on the same side of each discontinuity of ζ_i, and its complement. On the first set the difference is controlled by the Lipschitz constants of the

continuous parts, and its integral is at most $L_1 D$, using the integral bounds $\bar{z}$ and q in place of suprema. The second set has measure at most $2D$ for each chamber, because each of the two discontinuity phases lies between $s+\psi_i(s)$ and $s+\psi_i(k)$ only for s in an interval of length at most $|\psi_i(s)-\psi_i(k)| \le D$, and on it the difference is at most KJ. Thus the first and last terms contribute at most $E = D(L_1+\Lambda+2KJ) = \epsilon' C$ per window, complete or not. For the unfinished window, let $u = \tau - N$ and $g(s) = F(s,\hat{\psi}) - \bar{F}(\psi(N))$, which is one-periodic with zero mean and satisfies $|g| \le 2M$ componentwise, since $|F| \le M$ and so $|\bar{F}| \le M$. Then $\int_N^{N+u} g = -\int_{N+u}^{N+1} g$, so $\left|\int_N^{N+u} g\right| \le 2M \min(u, 1-u) \le M$. This is a contribution of order one to the integral (a direct bound of g by $2M$ gives $2M$ on the partial window; the factor M follows from the zero mean over the complete period, which lets the shorter of the partial interval and its complement be used), hence of order ϵ' to $\psi - y$, and it is the reason the bound is first order in ϵ' at a fixed time: it is the oscillating part of the solution, the displacement of ψ from its window average, and it cannot be cancelled by a window average. Writing $e(\tau) = \|\psi(\tau) - y(\tau)\|_\infty$, we have $\psi(\tau) - y(\tau) = \epsilon' \int_0^\tau [F(s,\psi_s) - \bar{F}(y(s))] ds$, so for every τ

$$e(\tau) \le \epsilon'^2 (\tau+1) C + \epsilon' M + \epsilon' \Lambda \int_0^\tau e(s)\, ds, \tag{34.8}$$

and Gronwall's inequality with the nondecreasing forcing $b(\tau) = \epsilon'^2(\tau+1)C + \epsilon' M$ gives $e(\tau) \le b(\tau) e^{\epsilon'\Lambda\tau}$. At an integer $\tau = N$ the unfinished window is empty, so the forcing at N is $\epsilon'^2 NC$ alone, while the integral term is bounded using $e(s) \le b(N) e^{\epsilon'\Lambda s}$ for $s \le N$, which gives $\int_0^N e \le b(N)(e^{\epsilon'\Lambda N} - 1)/(\epsilon'\Lambda)$ and the stroboscopic estimate. At fixed N the bracket $e^{\epsilon'\Lambda N} - 1$ is of order ϵ', so the whole is of order ϵ'^2 at fixed N and not at $N \sim 1/\epsilon'$. □

Correction of an earlier version. A first version of this theorem summed the complete-window contributions and applied the cancellation of the frozen oscillatory term to an arbitrary terminal time, which gave $\epsilon'^2(\tau+1)Ce^{\epsilon'\Lambda\tau}$ for every τ. That is not valid at times that end inside a window, where the frozen term has not completed a period. The statement above carries the order-one contribution of the unfinished window, and the sharper estimate of order ϵ'^2 is retained only at fixed integer times, not over the horizon $1/\epsilon'$ of the averaged flow. For comparisons of instantaneous offsets the first-order term is the displacement correction W of the offset paragraph below, which is the same oscillatory displacement and which the offset comparison included; the theorem itself bounds the displacement only through M.

Proposition 34.1 (Persistence of a locked state)**.** *Let the averaged relative-phase flow $w' = \epsilon' R\bar{F}(w)$, with w the differences of the components of y, have an equilibrium w^* and satisfy $\|w(\tau) - w^*\| \le K_c e^{-\epsilon'\rho\tau}\|w(0) - w^*\|$ whenever $\|w(0) - w^*\| \le \rho_0$. Put $L = \ln(2K_c)/\rho$ and $\Delta = 2\epsilon'\big[M + (L+\epsilon')C\big]e^{\Lambda L}$. If $\Delta \le \rho_0/4$ and the exact relative phases start within $\rho_0/2$ of w^*, then*

$$\limsup_{\tau\to\infty} \|w_{\text{exact}}(\tau) - w^*\| \le (1 + 2K_c)\Delta. \tag{34.9}$$

Proof. Let $r_k = \|w_{\text{exact}}(kL/\epsilon') - w^*\|$. On each interval of length L/ϵ' compare the exact solution with the averaged one started from the exact state. Theorem 34.2, applied from the start of the interval (the windows are aligned with its start, and the periodicity of the forcing makes the restart legitimate) and to differences with a factor 2, bounds their distance by Δ, and the averaged solution has contracted its distance to w^* to at most $K_c e^{-\rho L} r_k = r_k/2$

whenever $r_k \le \rho_0$. Hence $r_{k+1} \le r_k/2 + \Delta$ and $r_k \le 2^{-k} r_0 + 2\Delta \le \rho_0$ for all k, so $\limsup r_k \le 2\Delta$. Inside an interval, the distance is at most $\Delta + K_c r_k$, which gives the stated limit. □

Value of the constants. For the networks of Chapters 32 and 33 the constants computed from the cycles are as follows, with $\tilde{\delta} = 0.1$, $K = 2$ and the relaxation rates 0.156 and 0.234 per unit time of the averaged models.

Cycle	ϵ'	M	$\bar{d}$	γ	Λ	q	L_1	J	C	ρ
A	0.054	1.76	0.51	1.00	4.0	1.12	8.0	1.10	43	2.6
B, retuned	0.163	3.00	0.50	1.00	4.0	1.99	13.1	1.54	104	2.5

Proposition 34.1 is conditional, and its hypotheses are not established at these couplings. It needs a contraction neighborhood of the averaged relative-phase flow with certified constants K_c and ρ_0, whereas the column ρ above is the linear relaxation rate of the averaged model at its equilibrium and no neighborhood or K_c has been certified. It also needs $\Delta \le \rho_0/4$. The lifted flow is invariant under integer translations of the phases, so translates of w^* are equilibria and ρ_0 cannot exceed one cycle. Evaluating the expression with the favorable choice $K_c = 1$ and the rounded constants of the table gives $\Delta \approx 4.9$ cycles for cycle A and $\Delta \approx 48$ for cycle B, and $(1 + 2K_c)\Delta \approx 15$ and 145 cycles. Both exceed any admissible $\rho_0/4$ by more than an order of magnitude, so the hypothesis $\Delta \le \rho_0/4$ fails, and the evaluated numbers are values of a formula, not ultimate bounds on the locked offsets at these parameters. The proposition says that persistence of a locked state can be established when Δ is below the contraction radius. Since Δ is proportional to ϵ' at leading order for fixed L, through the term $\epsilon' M$ and the term $\epsilon'(L + \epsilon')C$, this needs ϵ' of order 10^{-3} or less: $\Delta/\epsilon' \approx 90$ for cycle A and 300 for cycle B. For comparison, the observed differences between the averaged and hybrid locked offsets are 0.0002 cycle in cycle A and 0.0012 in cycle B for the loaded networks, and 4×10^{-5} in the unloaded-equivalent network, so the formula exceeds them by four to five orders of magnitude. The weakness is the window argument, which bounds each window by a worst case over the displacement allowed by M, whereas the actual displacement in a locked state is a small fraction of it. What the theorem does establish is the form of the error: first order in ϵ' at a fixed time, with a constant that grows with the frequency spread through M, with the delay through $\bar{d}$, and with the discontinuity of the phase response through J. Those three dependences are the ones tested next.

34.3 Loaded Chambers

A bladder load changes the structure of the problem, and the first question is whether pressure and valve state still determine the evolution. They do not. With a bladder of pressure b_i, linear leak conductance γ_i and vent ratio ν, chamber i obeys

$$\dot{p}_i = \frac{g_{s_i} - p_i}{a_{s_i}} - \gamma_i (p_i - b_i) + q_i, \qquad \dot{b}_i = \gamma_i (p_i - b_i) - \nu \gamma_i b_i, \tag{34.10}$$

so the state is (p_i, b_i, s_i) and b_i is an independent continuous variable. Two chambers with the same (p, s) and different b evolve differently, so the single-pressure phase of Theorem 34.1

is no longer a coordinate on the state space, and a phase-only equation needs either a further assumption or a controlled approximation. The structure can nevertheless be written exactly.

In each mode the uncoupled flow is affine in (p, b) and is solved in closed form. Let Σ be the section $\{p = p_l,\ s = 0\}$, the state at a valve closing, parametrized by the bladder pressure η. Every state with $p \in [p_l, p_h]$ visited in the weak-coupling regime lies on the uncoupled flow line through a unique point $(p_l, \eta) \in \Sigma$ at a unique elapsed time θ, and (θ, η) are flow-box coordinates. Uncoupled, $\dot\theta = 1$ and η is constant between closings; at a closing, $\theta \mapsto \theta - T(\eta)$ and $\eta \mapsto M(\eta)$, where T is the time and M the bladder pressure at the next return to Σ. Both are analytic in η by the implicit function theorem when the switchings are transversal.

Theorem 34.3 (Exact equations of a loaded chamber)**.** *Assume that the crossings of p_h and p_l are transversal and that the coordinates (θ, η) are defined on the set visited. Then, with q_i the gate inflow, the loaded chamber obeys exactly*

$$\dot\theta = 1 + Z(\theta, \eta)\, q, \qquad \dot\eta = I(\theta, \eta)\, q, \qquad Z = \frac{\partial \theta}{\partial p}, \quad I = \frac{\partial \eta}{\partial p}, \tag{34.11}$$

with the reset $(\theta, \eta) \mapsto (\theta - T(\eta), M(\eta))$ when $\theta = T(\eta)$, and the pressure is $p = P(\theta, \eta)$. If M has a fixed point η^* with multiplier $\lambda = M'(\eta^*) \in (0, 1)$, then freezing $\eta = \eta^*$ in the first equation gives a single-pressure-type phase equation with the loaded cycle's own period $T(\eta^*)$, waveform $P(\theta, \eta^*)$ and response $Z(\theta, \eta^*)$. This is an approximation and not an exact subsystem: since $\dot\eta = I(\theta, \eta)q$, the set $\eta = \eta^*$ is invariant under coupling only if $I(\theta, \eta^*)q \equiv 0$, which fails whenever the gate inflow q is nonzero at phases where $I \neq 0$.

Proof. (θ, η) is a smooth function of (p, b) in each mode, and the gate inflow enters only the equation for p. By the chain rule $\dot\theta = \partial_p\theta\, \dot p + \partial_b\theta\, \dot b$, and the sum of the terms not containing q is the value 1 that θ has along the uncoupled flow by construction, leaving $\partial_p\theta\, q$; likewise $\dot\eta = \partial_p\eta\, q$ because η is constant along the uncoupled flow within a cycle. The coordinates are continuous across the mode switch at p_h because they follow the flow through it. The reset is the return to Σ. □

The equation is exact, within the stated domain, and has two variables per chamber. The phase-only model of a loaded chamber is obtained by freezing $\eta = \eta^*$, with its own period, waveform and response, which are the quantities that enter the first-order model with the loads as constant frequency shifts. It is an approximation, because the coupled chamber does not remain at η^*; what it leaves out is the motion of the slow variable. The sensitivities that govern it were measured for the free-leg chamber of cycle B, with $\gamma/C = 0.0051$ and the nominal supply, by solving the return map and by kicking the pressure at twenty phases in the hybrid simulation.

Quantity	Value
Return multiplier λ of the bladder pressure	0.98644 (against $e^{-(1+\nu)\gamma T} = 0.98645$)
Relaxation rate $\kappa = \ln\lambda/T$, time constant $1/\lvert\kappa\rvert$	-0.0077, 130 time units
Period sensitivity $d\ln T/d\eta$ per unit bladder pressure	-0.0018
Response sensitivity $\max\lvert\partial Z/\partial\eta\rvert/\max\lvert Z\rvert$	0.0018
Isostable response $\max\lvert I\rvert$, change of closing bladder pressure per unit kick	0.0052

With $|q| \le \sum_j \epsilon_{ij} a$ a priori, the linearization of the slow variable about η^*, with forcing $\dot{\eta} = Iq$ and relaxation at $|\kappa|$, suggests a deviation of the order of $\sup |I|\,|q|/|\kappa| = 0.0052 \times 0.74 \times 130 \approx 0.5$ pressure units for the retuned coupling, which would shift the frequency of the chamber by 9×10^{-4}. This is a local linear estimate, not a bound: the multiplier and the sensitivities were measured at the single fixed point η^*, and a nonlinear stability argument over the region visited has not been made. In the hybrid simulation of the locked network the closing bladder pressures differed from their uncoupled values by at most 0.0005, 0.003 and 0.012 pressure units at extra detuning 0, +0.010 and +0.020, and by at most 0.013 at −0.014, so the frequency effect is below 3×10^{-5} and the effect on the response below 3×10^{-5} in relative terms, about 40 times smaller than the a priori bound and about 10^{-3} of the margins of locking. This is consistent with the measurement that freezing the bladder changes the limits of locking by at most 0.0002. The slow variable relaxes at a rate 0.004 to 0.03, much smaller than the coupling rates 0.05 to 0.2, so on the time scale of phase locking it is a parameter drifting at order ϵ, and the system to be approximated by averaging is the phase model with parameters η_i, one per chamber, not the phase model alone. Whether Theorem 34.2 extends to the reset form, with η as an extra slow variable, is not proved here; its constants would acquire the sensitivities above, which are of relative size 10^{-3}.

34.4 Diagnosis of the Cycle B Discrepancy

In cycle B the limits of locking predicted by the first-order averaged model differ from the hybrid ones for the loaded network (Chapter 37). The size of the difference depends on which prediction is compared with which hybrid run, and every percentage in this Part is the difference between a predicted limit and the hybrid limit divided by the *predicted* limit, taken at each end of the range separately. Three predictions appear: P1 keeps the nominal interaction functions and changes only the frequency of the detuned chamber (this is the averaged model of Chapter 37); P2 recomputes the interaction functions for the detuned chambers with the loads entered as frequency shifts; and P2′ recomputes them for all three chambers when the loads are replaced by equivalent supply shifts. The following values are for the retuned network at $\epsilon T_{SF} = 0.163$, with the hybrid limits $[-0.0140, +0.0207]$ for dynamic bladders and $[-0.0140, +0.0214]$ for the equivalent supply shifts.

Prediction	Compared with	Lower limit, percent	Upper limit, percent
P1, $[-0.0173, +0.0258]$	dynamic bladders	19	20
P2, $[-0.0164, +0.0250]$	dynamic bladders	15	17
P2′, $[-0.0163, +0.0248]$	dynamic bladders	14	17
P2′, $[-0.0163, +0.0248]$	equivalent supply shifts	14	14

Where the text says about twenty percent it refers to P1 against the dynamic hybrid, and the figure of fourteen percent refers to P2′ against the equivalent-supply hybrid at the lower limit. Candidate explanations were tested in turn, each against the same event-aware hybrid simulation. Unless stated, the cycle is B with the retuned network, the delay is $T/2$ and the detuning is applied to the free-leg chamber F by its supply pressure.

Offsets are not the discrepancy. The offsets of the locked states were compared first, because the limits of locking depend on them. With the hybrid offset taken as the lag between

valve closings divided by the *nominal* period, the hybrid and averaged offsets of Chapter 32 differed by $0.0014T$ and $0.0028T$ in cycle A. That comparison was not between the same quantities. The averaged model gives phase offsets $x_i - x_S$ in cycles of the chambers, and the lag divided by the nominal period differs from it by the relative frequency difference times the offset, which is 0.0025×0.5 for the free leg. The two are compared properly by reading the phase θ_i/T_i of each chamber from its pressure at the closing of S. The averaged field also has to be corrected for the fact that the instantaneous phase is not the mean phase: at the closing of S the offset is $y_i - y_S + W_i - W_S$, with W the zero-mean primitive of $F - \bar{F}$ evaluated at the closing phase. The comparison below uses the equivalent-supply-shift network (no bladders), for which the phase is read exactly from the pressures.

Case	Averaged (u_F, u_A)	With W	Hybrid (pressures)	Largest difference
A, $\epsilon T_{SF} = 0.054$	0.50624, 0.01562	0.50625, 0.01564	0.50625, 0.01564	0.00000
A, $\epsilon T_{SF} = 0.109$	0.50140, 0.00791	0.50141, 0.00792	0.50141, 0.00792	0.00000
B, $\epsilon T_{SF} = 0.082$	0.51838, 0.02638	0.51872, 0.02674	0.51888, 0.02691	0.00017
B, $\epsilon T_{SF} = 0.163$	0.50851, 0.01203	0.50867, 0.01219	0.50870, 0.01223	0.00004

With the offset defined consistently the averaged model reproduces the locked offsets to 2×10^{-4} or better in both cycles, and the W correction, of size up to 0.0004, is resolved in cycle B. The same holds for a pair without delay: the first-order averaged phase offset of two diffusively coupled chambers detuned by $0.12\epsilon T$ differs from the hybrid phase read from the pressures by 0.0001 to 0.0007 for ϵT from 0.02 to 0.16. The offsets of the locked states therefore carry no discrepancy that needs explaining, and an earlier attribution of an offset discrepancy to first-order averaging error was an artefact of the comparison. The limits of locking are a different quantity, determined by the folds of the dynamics and not by the offsets, and their discrepancy in cycle B stands.

A detuned chamber is a different chamber. Changing the supply of F by an amount that shifts its rate by δ changes its period, but also the fraction of the cycle spent charging, the waveform in normalized phase and the phase response. The table compares the detuned chamber with the nominal one, with phase expressed in cycles from the valve closing, δ equal to the rate change, and the response normalized to the pressure amplitude. The response was computed analytically and measured in the hybrid simulation by kicking the pressure at forty phases; the two agree to 0.04 percent of the peak for the unloaded chamber.

δ	$\Delta\sigma$	T_F	t_c/T	waveform change	response change
−0.0140	−0.063	1.7928	0.8668	2.1%	0.6%
−0.0100	−0.045	1.7856	0.8662	1.5%	0.4%
0	0	1.7677	0.8649	0	0
+0.0100	+0.045	1.7502	0.8635	1.5%	0.4%
+0.0200	+0.090	1.7330	0.8622	3.0%	0.8%
+0.0207	+0.093	1.7319	0.8621	3.1%	0.8%

The waveform change is the largest difference of the pressure in normalized phase from that of the nominal chamber, as a percentage of the amplitude $p_h - p_l$; the response change is the relative root-mean-square difference of the normalized response away from the switching

phases. Most of the change of angular frequency is the intended one. The rest is a shift of the switching phase by up to 0.003 cycle, a waveform change of a few percent and a response change under one percent. A bladder load at the periodic steady state lengthens the period by 0.36 percent, changes the waveform by up to 0.57 percent of the amplitude and the response by 0.7 percent, and the measured response of the loaded chamber differs from the analytic one by 1.1 percent. These are small, and the question is whether they matter for the limit.

Frequency only against recomputed interaction. Two predictions of the limits of locking were compared with the hybrid simulation. The first (P1) keeps the nominal interaction functions and changes only the frequency of the detuned chamber. The second (P2) recomputes all interaction functions from the changed chambers, using the chamber-specific Γ_{ij} of this chapter, with the loads entered as frequency shifts. The brackets of the hybrid simulation are those of Chapter 37.

Case	ϵT_{SF}	P1, frequency only	P2, recomputed	Hybrid
A, loaded	0.054	$-0.0156,\ +0.0199$	$-0.0153,\ +0.0197$	$-0.0151,\ +0.0198$
A, loaded	0.109	$-0.0304,\ +0.0403$	$-0.0296,\ +0.0400$	$-0.0286,\ +0.0406$
A, loaded	0.163	$-0.0450,\ +0.0600$	$-0.0432,\ +0.0600$	$-0.0409,\ +0.0616$
B, loaded	0.082	$-0.0094,\ +0.0100$	$-0.0091,\ +0.0098$	$-0.0085,\ +0.0076$
B, loaded	0.109	$-0.0121,\ +0.0156$	$-0.0117,\ +0.0153$	$-0.0105,\ +0.0124$
B, loaded	0.163	$-0.0173,\ +0.0258$	$-0.0164,\ +0.0250$	$-0.0140,\ +0.0207$
B, loaded	0.217	$-0.0224,\ +0.0353$	$-0.0208,\ +0.0340$	$-0.0167,\ +0.0278$
B, no loads	0.082	$-0.0086,\ +0.0140$	$-0.0084,\ +0.0139$	$-0.0078,\ +0.0136$
B, no loads	0.163	$-0.0168,\ +0.0281$	$-0.0159,\ +0.0274$	$-0.0136,\ +0.0262$

Recomputing the interaction moves the predicted limits by 0.0001 to 0.0018, a few percent, in the direction of the hybrid simulation at the lower limit and only slightly at the upper one. The assumption that detuning changes the frequency only is therefore a small error in both cycles, and recomputing it does not remove the cycle B discrepancy: for the loaded network at $\epsilon T_{SF} = 0.163$ the remaining gaps are 0.0024 and 0.0043.

Load mechanism. Three hybrid variants were compared for the loaded cycle B network: the bladders evolving dynamically, the bladders frozen at their periodic mean pressure, and no bladders with the load frequency shifts reproduced by equivalent supply changes of S, F and A. For the last, P2 was recomputed with the correspondingly detuned cycles for all three chambers (P2′).

ϵT_{SF}	Dynamic bladders	Frozen bladders	Equivalent supply shifts	P2′
0.082	$-0.0085,\ +0.0077$	$-0.0085,\ +0.0076$	$-0.0085,\ +0.0081$	$-0.0091,\ +0.0095$
0.163	$-0.0140,\ +0.0207$	$-0.0140,\ +0.0205$	$-0.0140,\ +0.0214$	$-0.0163,\ +0.0248$

Freezing the bladder changes the limits by at most 0.0002, so the bladder dynamics do not account for the discrepancy. Replacing the leak by an equivalent frequency shift changes the upper limit by 0.0004 to 0.0007. The discrepancy therefore remains with no bladders at all, and it is a discrepancy between the averaged equations and the hybrid chambers with unequal natural frequencies. Single-chamber load shifts in the hybrid simulation agree with the first-order formula to 0.00003.

Scaling with coupling and frequency spread. The bound of Theorem 34.2 predicts an error that is first order in ϵ' at a fixed ratio $\tilde{\delta} = \delta/\epsilon'$ of frequency spread to coupling, and a constant that grows with $\tilde{\delta}$. The equivalent-shift model was therefore run with the shifts scaled by a factor κ (so $\kappa = 1$ is the load spread of the network) and the coupling by a factor m (so $m = 3$ is the retuned network), keeping κ/m fixed in one family and κ fixed in another. The entries are the ratio of the hybrid limit to P2′.

m	κ	ϵT_{SF}	P2′ limits	Hybrid limits	Ratios, lower and upper
0.75	0	0.041	−0.0043, +0.0070	−0.0042, +0.0069	0.97, 0.99
1.5	0	0.082	−0.0084, +0.0139	−0.0078, +0.0136	0.93, 0.98
3.0	0	0.163	−0.0159, +0.0274	−0.0136, +0.0262	0.85, 0.96
0.75	0.25	0.041	−0.0045, +0.0065	−0.0044, +0.0063	0.97, 0.98
1.5	0.5	0.082	−0.0088, +0.0127	−0.0081, +0.0121	0.93, 0.95
3.0	1.0	0.163	−0.0163, +0.0248	−0.0140, +0.0214	0.86, 0.86
0.75	1.0	0.041	−0.0052, +0.0004	−0.0050, +0.0000	0.98, 0.10
1.5	1.0	0.082	−0.0091, +0.0095	−0.0085, +0.0081	0.93, 0.85
3.0	1.0	0.163	−0.0163, +0.0248	−0.0140, +0.0214	0.86, 0.86

In the first two families the discrepancy of P2′ against the equivalent-supply hybrid falls with the coupling, from about 14 percent to 3 percent in the family with scaled spread, roughly in proportion to ϵ', as an error of first order in ϵ' at fixed $\tilde{\delta}$ does. In the third family the spread is held fixed while the coupling is reduced, so $\tilde{\delta}$ grows; at $\epsilon T_{SF} = 0.041$ the spread consumes nearly the whole upper range and the limit is a small difference of large numbers, which makes the ratio 0.10 uninformative, and at 0.082 the upper limit of P2′ is 15 percent off the equivalent-supply hybrid. The loaded discrepancy that did not seem to fall with coupling was therefore a fixed absolute spread in a shrinking locking range, which is outside the regime in which the frequency spread is of the order of the coupling.

Second-order averaging. The next term of the averaging expansion was evaluated numerically. With $W(s, y)$ the mean-zero primitive in s of $F(s, y) - \bar{F}(y)$, the second-order averaged field is $\bar{F}(y) + \epsilon' F_2(y)$ with $F_2(y) = \frac{d}{dh}\big\langle F(s, y + hW(s, y))\big\rangle_s\big|_{h=0}$, computed by central differences so that the discontinuity of ζ_i is handled by the averaging. For a delayed edge the partner fluctuation is evaluated at the delayed phase, $W_j(s - d_{ij})$, and the drift of the partner during the delay enters through the collective rate. This is the standard second-order term, not proved here for a discontinuous right-hand side. For the diffusive pair of cycle A the term is $F_2 \approx 5 \times 10^{-6}$ against a first-order field of 8.5×10^{-3}, too small to matter.

m	κ	P2′	Second order	Hybrid
0.75	0.25	−0.0045, +0.0063	−0.0045, +0.0064	−0.0044, +0.0063
1.5	0.5	−0.0088, +0.0127	−0.0088, +0.0110	−0.0081, +0.0121
3.0	1.0	−0.0163, +0.0248	−0.0163, +0.0246	−0.0140, +0.0214
1.5	1.0	−0.0091, +0.0095	not reliable, +0.0094	−0.0085, +0.0081
1.5	0	−0.0084, +0.0139	−0.0084, +0.0138	−0.0078, +0.0136
3.0	0	−0.0159, +0.0274	not reliable, +0.0270	−0.0136, +0.0262

The second-order term does not account for the discrepancy. It changes the upper limit by 0.0001 to 0.0004 in four cases, moves it away from the hybrid value in none of them

by more than that, and in one case, $m = 1.5$ with scaled spread, it lowers the upper limit by 0.0017 to a value 0.0011 below the hybrid one, overshooting. It leaves the lower limit unchanged where the continuation was reliable. The entries marked not reliable are those where the continuation of the root failed. The residual discrepancy of the limits is therefore not accounted for by the first-order model, by recomputing its functions for the detuned chambers, by the bladder dynamics, by the offsets, or by the second-order averaging term as computed here; these tests show that each correction is insufficient, and they do not show that no such mechanism contributes.

Phase drift and unsettled runs. The brackets of Chapter 37 were obtained from runs of 3000 time units, and near a limit a slow settling can be mistaken for a drift. The upper limit of the retuned loaded network was examined with runs of 6000 time units. At extra detuning +0.0200 the offset of F moved by 0.17 cycle in the first 1000 time units and not afterward. At +0.0207 it moved by 0.19 cycle in the first 1000, by 2.48 cycles in the next 1000, with two slips, and not afterward: the run locks, after a settling of about 2000 time units. At +0.0212 it slipped continuously, 51 slips in 6000 time units, with the offset moving by 8 to 9 cycles in every window of 1000. Drift is therefore demonstrated at +0.0212 and above, the run at +0.0207 shows that lock can appear only after a long transient, and the upper limit lies between 0.0207 and 0.0212, in agreement with the bracket.

34.5 What Is Established and What the Bound Must Cover

Two statements about exactness need to be kept apart. The unaveraged phase equation (34.2) of a single-pressure chamber is exact, and so is the two-variable equation of Theorem 34.3 for a loaded chamber. The averaged interaction equation of Chapter 29 is an approximation to either. For a single-pressure chamber its error is the whole distance between the phase model and the chamber; for a loaded chamber the phase model also freezes the slow variable at its fixed point, a further approximation, and the distance has both parts.

The reduced state space for that error is the phase model with chamber-specific functions and delays: one phase per chamber without loads, and one phase plus one slow variable η_i per chamber with loads. The slow variable relaxes at 0.004 to 0.03, against coupling rates of 0.05 to 0.2, and its measured effect on the frequency and response is of relative size 10^{-5}, so it can be carried as a parameter and its dynamics neglected at the level of the observations; no theorem here covers that neglect. A bound built on the single-pressure model alone covers the loaded network only up to that measured effect.

Of the original candidates, a heterogeneous-chamber reduction is the correct averaged model and changes the limits of locking by a few percent. Explicit load-state dynamics are not needed at the level tested. A sharper averaging estimate was sought through the second-order term, and the term as computed does not account for the discrepancy. The locked offsets, once measured as phases, agree with the averaged model to 2×10^{-4} or better, so the discrepancy is confined to the limits of locking, which are determined by the folds of the dynamics. For a delayed pair the hybrid limits agree with the first-order ones computed from the heterogeneous tables, with the collective-rate delay phase, to 0, 3 and 8 percent of the predicted lower limit at $\epsilon T_{SF} = 0.04$, 0.08 and 0.16 in cycle B, growing with the coupling in the way an averaging error would, although that trend alone does not identify its cause; the

loaded three-chamber network is 14 to 17 percent inside the recomputed prediction and 19 to 20 percent inside the frequency-only one at the same coupling (percentages of the predicted limit, as defined at the start of this section). The cause of the larger network discrepancy is unresolved. The offsets, the bladders, the waveform change of the detuned chamber and the second-order term as computed were tested and are each insufficient to account for it, and it grows with the coupling at fixed ratio of spread to coupling, which is consistent with an averaging error without establishing one. Agreement of the locked offsets, which are interior solutions, does not establish agreement of the boundaries of locking, which are determined by the folds. The lower limit remains unexplained and is treated as a separate open diagnostic: a correction of higher order is the leading candidate, and it has not been tested.

The sharper bound was not derived. What the chapter provides is a rigorous bound with constants too large to be informative, a measurement of where its slack lies (the first-order averaged offsets are accurate to 2×10^{-4} in cycle A where the formula of the persistence proposition evaluates to about 15 cycles, a value that is itself not an established bound because the proposition's hypotheses fail), and a specification of the model a sharper bound must treat: the exact phase equation with chamber-specific functions, discontinuous phase response, delays, and frequency spread of the order of the coupling. The route to a sharper bound is a window-map expansion to second order with an explicit remainder, in which the discontinuity of ζ_i contributes through the measure of the set where a trajectory and its frozen comparison lie on different sides of a switching phase; its terms were evaluated numerically above and not bounded.

What remains open is a sharp bound on the averaging error and on the limits of locking that is small enough at $\epsilon' \approx 0.05$ to 0.2 to be tested; the cause of the larger discrepancy of the loaded network; the extension of the averaging theorem to the reset form with the slow variable; and the case of a shared supply, which couples the chambers through a state-dependent term not covered by Assumption W.

Chapter 35

Didactic Tasks as Coordination Sequences

The author holds that robots should have a didactic purpose, and the two recorded examples can be written as sequences of the operations of Part III. The theory addresses the timing of coordinated motion and not the content of a task, so each example is stated in terms of what the theory does and does not cover.

35.1 Applying Paper Sheets

A robot that prepares itself by applying sheets of paper one by one must hold a stance while one arm group reaches, applies a sheet and returns. In the language of Part III the stance group g, consisting of the legs and trunk, is retained, and the arm group h is released, acts independently and is recaptured. A cycle consists of maintained coordination, release of the arm group, independent action during which the arm group follows its own internal pattern, and recapture. The theorems give three conditions for the cycle to work. The stance group must stay cohesive throughout, which Theorem 13.3 gives for every schedule of the external coupling when the floor margins hold. The release must be strong enough, in the sense of its inequality, that the arm is not dragged by the stance rhythm. And the recapture must begin with the relative phase outside the excluded band, so that recapture is guaranteed. Because the relative phase drifts during release, a schedule that releases for a fixed duration may end inside the band, and a practical remedy is to wait until the phase leaves the band or to apply a brief detuning pulse before restoring coupling. Contact with the sheet is a mechanical event and belongs to the hybrid-consistency obligation, which has been proved only for the cycling chamber.

35.2 Standing on One Leg While Dressing

A robot that jokes with a child while standing on one leg and pulling on trousers has at least three groups: the stance leg with the trunk, the free leg, and the arms. Chapter 25 treats this grouping. The three-group theorem there establishes that every group keeps its internal coordination while any group is released, and the pairwise-exclusive schedule gives release and

recapture for one pair at a time. Simultaneous recapture among the three groups is proved in Theorem 25.3, but only under consistent group targets and a margin condition on the weakest cross weight; the frustrated case is open, and in the first embodiment purely diffusive gates between identical chambers do not supply consistent nonzero offsets without designed delays. Two further limits should be stated plainly. Balance on one leg is a mechanical stability problem, in which the projected center of mass must remain over the support, and it is not addressed by the oscillator theory, which provides coordinated actuation signals and does not guarantee that they keep the body upright. The jokes are behavior at a level above motion timing, and the theory contributes only the sequencing of the motions that accompany them.

35.3 What the Examples Test

Both tasks require a retained group to stay coordinated while another group acts independently and then rejoins. They are therefore tests of the central claim of the monograph, and a body that can perform them with measured cross coherence falling during release and recovering after recapture would support the theory in the sense of the discriminating observations of Chapter 10.

Chapter 36

Validation Program

The claims of the monograph are meant to be testable, and this chapter collects the tests. Several are adapted from the empirical program of the locomotion essay [25], which proposes tests of compactness of the reduced description, of the contact reset, and of balance recovery. They are stated for the endomarionette and for biological gait where the essay's data are applicable, and no test has been run.

36.1 Estimating Phase and Coherence

The essay's estimator takes Hilbert-transform analytic signals of joint angles and averages $e^{\mathrm{i}\hat{\theta}_i}$ over the joints. It is an estimator of kinematic phase and not a direct measurement of oscillator phase, and it averages unadjusted phases, so an orderly antiphase movement would read as low coherence. The estimator used here is pattern adjusted: $\hat{r}_a(t) = \left| n_a^{-1} \sum_{i \in a} e^{\mathrm{i}(\hat{\theta}_i(t) - \hat{\beta}_i)} \right|$, with the target offsets $\hat{\beta}_i$ taken as the circular means of $\hat{\theta}_i - \hat{\theta}_{\mathrm{ref}}$ over baseline windows of steady movement. The offsets must be estimated on windows separate from those in which coherence is evaluated, since using the same data for both inflates $\hat{r}_a$ toward one.

For the pneumatic body the cycle timestamps are observable directly, because each cycling chamber has a valve event, but the phase between events is not measured and needs a stated interpolation or model. Two options are available. Linear interpolation, $\hat{\theta}(t) = 2\pi(t - t_k)/(t_{k+1} - t_k)$ between consecutive closing events, assumes a constant phase rate within the cycle, which holds for the unperturbed oscillator and fails when perturbations act inside the cycle, since the phase response of Chapter 29 is not uniform; the interpolation is then correct at the events and approximate between them. The second option uses the pressure itself: if the chamber pressure is sensed, the unperturbed waveform is invertible on each branch, giving $\hat{\theta} = -\tau_c \ln \frac{p_s - p}{p_s - p_l}$ while charging and $\hat{\theta} = t_c - \tau_o \ln \frac{p - p_o}{p_h - p_o}$ while discharging, which is a model-based phase between events and inherits the model's idealizations. For cord-driven joints the phase of the winding channel is the corresponding measurement.

36.2 Tests

Cohesion bound. Measure the within-group spread $D_a(t)$ and compute γ_a from independently measured quantities: the algebraic connectivity of the internal coupling graph, the

dispersion Ω_a of the natural frequencies, and the external coupling degrees. The theorems predict $D_a(t) \le \gamma_a$ when the initial condition lies in the invariant set. An observed violation falsifies the hypotheses, for instance sinusoidal coupling, or the measured parameters.

Release boundary. Sweep the mean frequency mismatch $\Delta\bar{\omega}$ by a drive, or sweep the cross coupling, and locate the threshold between locking and drift. The theorems predict that the threshold lies between $\bar{N}(1 - \gamma_g - \gamma_h)$ and $\bar{N}(1 + \gamma_g + \gamma_h)$ in units of $|\Delta\bar{\omega}|$, with locking guaranteed below the lower end and drift guaranteed above the upper end. A threshold outside the interval falsifies the reduction or the measured coupling.

Recapture. Restore the coupling with the relative phase set at chosen values, including values inside the excluded band, and record whether and how fast u enters the arc. The theorems guarantee capture only outside the band. The band contains the unstable equilibrium, and the simulations of Chapter 16 are not substitutes for these trials.

Transition mechanisms. Record c_{gh} and Φ during pattern changes and compare with the four signatures of Chapter 10.

Compactness. The essay proposes that the description length of a trajectory under the amplitwist chain should not exceed that under a direct joint-angle description, for steady gaits. It fails if the chain needs more parameters or a deeper hierarchy than the direct description. For the endomarionette the chain description is the frequencies, the coupling structure, the offsets and the schedule of the control, and the direct description is the joint trajectories.

Contact reset. The essay proposes computing the ratio of the summary just after and just before each foot strike and checking that it is approximately constant across steps of the same gait. By Proposition 18.1 the further requirement for the endomarionette is that the reset keep the spread inside its bound.

Balance recovery. The essay conjectures that coherence declines measurably before visible imbalance and then recovers, with the phase relocking. It proposes comparing the pattern-adjusted coherence between perturbation onset and the first instability crossing against matched unperturbed windows, with a paired comparison, and reanalyzing existing perturbation datasets before collecting new data. It does not specify a significance threshold or a sample size. The cohesion theorem gives a supporting mechanism, since an increase in external disturbance enlarges the bound γ_a, but it does not by itself predict a precursor of instability, and the conjecture remains empirical.

Levels of evidence. A test at one level does not stand for a test at the next. In a hardware precedent a closed-loop link between an artificial rhythm network and biological cells modified the network while the robot's movement remained disorganized, and the authors did not claim the behavior it was meant to support [1]. The tests above are therefore kept at separate levels: phase coordination among chambers, phase and amplitude at the drum coordinate,

joint motion, and behavior such as balance or recovery. Passing a test at one level supports a claim at that level, and a claim about a higher level needs its own test.

Chapter 37

Bench Protocol for Two Chambers

37.1 Status and Purpose

This chapter is a protocol and not evidence. No bench has been built and no measurement has been made, and every number given below as a prediction is a prediction of the model of this monograph, to be compared with a measurement that does not yet exist. The central output of the protocol is the measured interaction function $\hat{G}(\varphi)$ of two coupled chambers, its disagreement with the predicted function $G_{\mathrm{pred}}(\varphi)$ of Chapter 30, and the uncertainty of both. The protocol also yields the measured phase response, the locking range and the stable offsets. It is written so that a second person could execute it, and so that the result can contradict the theory.

37.2 Quantities

The model is $\dot{\varphi} = \delta + \epsilon G(\varphi)$ of Equation (30.1), with $\epsilon = g/C$. The protocol measures, for each chamber, the waveform $p(\theta)$, the thresholds, the time constants, the period T and its jitter, and the phase response curve $Z(\theta)$; for the pair, the natural rate difference δ; for the gate, the conductance g, its linearity and the capacitance C; and for coupled runs the phase-difference series $\varphi(t)$ from valve-event timestamps. One prediction is independent of the cycle parameters: for a symmetric diffusive gate between identical chambers $G'(0) = -2$, by the normalization $Z\dot{p} = 1$. It is parameter-free once the conductance and capacitance have been calibrated independently, and it depends on the scalings of Section 37.3.

37.3 Normalization of the Slope Prediction

The value $G'(0) = -2$ is a pure number only under the following scalings, and the bench has to reproduce them or the comparison is between a dimensional measurement and a dimensionless prediction. Time and phase: the phase θ of a chamber is measured in units of time and advances at unit rate along the cycle, so $\varphi = \theta_2 - \theta_1$ is a time difference, and $\dot{\varphi}$ is a time per time. Pressure: the phase response Z has the dimension of time per pressure and is normalized by $Z\dot{p} = 1$ on the cycle, so the product of Z and a pressure difference is a time, the unit of pressure cancels, and G has the dimension of time. Conductance and capacitance: the gate acts on chamber i by $\dot{p}_i = \epsilon\,(p_j - p_i)$, where $\epsilon = g/C$ is an inverse time,

g is the linear conductance, in mass flow per unit pressure difference, and C the capacitance, in mass stored per unit pressure, of the isothermal gas volume; for an ideal gas of specific gas constant R_s at temperature Θ in a rigid volume V, $C = V/(R_s\Theta)$. With these definitions both sides of $\dot\varphi = \delta + \epsilon G(\varphi)$ are dimensionless and $G'(0) = -2$ is dimensionless. Equivalently, a small phase difference decays at the rate 2ϵ per unit time, which is a fraction $2\epsilon T$ per cycle: $\epsilon T = 0.02$, 0.04 and 0.08 correspond to contractions of about 4, 8 and 16 percent per cycle at first order.

The slope is unchanged if phase is expressed as a fraction of the period or in radians, provided the rate $\dot\varphi$ and the phase use the same unit and ϵ is the inverse of the time unit in which the rate is written, because the same factor multiplies both axes. It is also unchanged by a change of pressure unit and by a change of the period. It is not unchanged by a mis-scaled ϵ: an error of the capacitance C by a factor enters the apparent slope by the same factor, which is why the reported statistic is the measured decay rate divided by the independently calibrated $\hat\epsilon = \hat g/\hat C$, compared with 2, and not a free fit. The prediction assumes identical chambers, a symmetric gate whose flow is linear in the pressure difference at the differences that occur near lock, a gas that is ideal and isothermal at the cycle frequency so that one capacitance describes the storage, instantaneous transport, and first order in ϵT. A gate that is orifice-like and not laminar has a flow that is not linear in small differences, and the slope at small φ then depends on the pressure difference and not only on g. A measured slope different from -2 therefore tests the derivation together with these conditions and with the calibration of g and C; Stage 0 measures each condition separately so that a discrepancy can be assigned.

Two properties of G affect how the slope is measured. First, G is not linear away from zero. For the cycle A of Chapter 33, $G(\varphi)/\varphi$ is -1.91 at $\varphi = 0.02T$, -1.73 at $0.06T$ and -1.55 at $0.1T$, and for cycle B it is -1.77, -1.36 and -0.99. Second, the departure is not a cubic term, and the leading correction can be derived. Changing variable in the second half of the definition gives $G(\varphi) = \langle Z(s)[p(s-\varphi) - p(s+\varphi)]\rangle$, which is odd. Away from the valve events p is smooth, so $p(s-\varphi) - p(s+\varphi) = -2\varphi\dot p(s) + O(\varphi^3)$, and with $Z\dot p = 1$ these regions give -2φ. Within $|\varphi|$ of an event, p is continuous with $\dot p$ jumping from a to b, and Z jumps from $1/a$ to $1/b$. If p is replaced on each side by its tangent line, the integrand differs from $-2\varphi\dot p(s)$ by $(a-b)(s-k+\varphi)$ before the event time k and by $(a-b)(s-k-\varphi)$ after it, and weighting by $1/a$ and $1/b$ and integrating over the two intervals of length $\varphi > 0$ gives $(a-b)(1/a - 1/b)\varphi^2/2 = -(a-b)^2\varphi^2/(2ab)$. At both events a and b have opposite signs, so this equals $(a-b)^2\varphi^2/(2|ab|)$, and by oddness

$$G(\varphi) = -2\varphi + b_G\,\varphi|\varphi| + O(\varphi^3), \qquad b_G = \frac{1}{2T}\sum_k \frac{(\Delta\dot p_k)^2}{|\dot p_k^- \dot p_k^+|}, \tag{37.1}$$

the sum running over the two valve events and the neglected terms coming from the curvature of p and Z on each side. For cycle A this gives $b_G = 5.16$ and for cycle B 6.67, against free fits over $|\varphi| \le 0.02T$ of 5.18 and 6.65. The term is therefore a consequence of a local expansion under stated assumptions, namely a continuous p that is smooth between two events, $Z = 1/\dot p$ on each side, and $|\varphi|$ small compared with the spacing of the events and with the time scales of curvature; it is not an arbitrary fitting choice, but the assumptions are those of the model and have not been checked for a real chamber. An odd cubic fit over $|\varphi| \le 0.05T$ returns a slope of -1.93 for cycle A and over $0.15T$ of -1.78.

The slope obtained depends on the window and on the form of the fit, as the following values show for the computed G. The columns are an odd cubic, $a_1\varphi + b\,\varphi|\varphi|$, the three-term form $a_1\varphi + b\,\varphi|\varphi| + a_3\varphi^3$ with b free, and the cubic with b fixed at b_G.

Cycle and window	Cubic	$\varphi\|\varphi\|$	Three terms	b fixed
A, $0.05T$	-1.927	-1.998	-2.000	-2.000
A, $0.1T$	-1.854	-1.993	-1.999	-2.000
A, $0.15T$	-1.781	-1.987	-1.999	-2.001
A, $0.2T$	-1.709	-1.979	-1.997	-2.003
B, $0.05T$	-1.816	-1.981	-1.998	-2.000
B, $0.1T$	-1.638	-1.941	-1.989	-2.008
B, $0.15T$	-1.470	-1.897	-1.979	-2.024
B, $0.2T$	-1.247	-1.728	-2.106	-1.985

With the term included the slope of the three-term fit is within 0.001 of -2 for windows up to $0.15T$ in cycle A, and within 0.002 at $0.05T$ and 0.011 at $0.1T$ in cycle B; with b fixed the deviation is at most 0.008 for windows up to $0.1T$ in both cycles. In cycle B the three-term fit degrades beyond $0.15T$, where structure that the local expansion does not contain enters. The protocol therefore takes the slope statistic to be the coefficient a_1 of the three-term fit over $|\varphi| \le 0.1T$, repeated with windows of $0.05T$ and $0.15T$ and with b fixed at b_G computed from the fitted waveform, and reports all four values; the comparator value is the coefficient of the same fit to the predicted G. These numbers are for noise-free G; with measurement noise a narrow window raises the variance of a_1, which is the trade the window choice makes and is propagated by the bootstrap of Section 37.9.

37.4 Apparatus

Element	Requirement	Reason
Two chambers	Rigid, equal volume, volume known to 1 percent, heat-conducting walls or insert to keep the gas close to isothermal	Capacitance C and the isothermal assumption
Valve with hysteresis	Opens at p_h and closes at p_l, switching time at most 1/50 of the shortest cycle timescale, thresholds stable to 0.5 percent over a session	Idealized switching of Chapter 2
Supply	Regulated, independent restriction per chamber, supply pressure logged and stable to 0.2 percent, with a configuration that makes shared-line coupling negligible and one that makes it measurable	Residual path through the supply
Discharge path	Fixed orifice per chamber	Time constant τ_o
Gate	Adjustable conductance, calibrated against flow and pressure difference, leak of the closed state measured; for a delayed gate a selectable tube with measured transport delay	g, its linearity, and τ
Pressure sensors	One per chamber, noise at most 0.2 percent of $p_h - p_l$, bandwidth at least 20 times the cycle frequency	Timestamps and waveforms
Acquisition	Simultaneous sampling at an interval not above $\min(\tau_o/100, T/500)$, one clock, threshold crossings located by linear interpolation	Timestamp resolution
Perturbation source	Solenoid or equivalent delivering a pulse of known volume, at most 5 percent of $p_h - p_l$, triggered at a chosen phase	Direct measurement of Z
Environment	Dry air, ambient temperature logged to 0.5 K, leak test by pressure decay with valves and gate closed	Isothermality and tightness

37.5 Stage 0: Calibration

The sensors are calibrated against a reference gauge at no fewer than eight pressures on rising and falling ramps, and their noise is taken from static records. The chamber volumes are measured by isothermal compression. The valve thresholds are measured by slow quasi-static ramps at three ramp rates, to check that they do not depend on the rate; a rate dependence is recorded and carried into the model. The gate conductance is measured as steady flow against pressure difference at no fewer than six pressures covering the operating range, and fitted to the linear model of Equation (21.1); the largest deviation from linearity, as a fraction of the flow, is reported, because the derivation assumes it is small. The transport delay of a tube is measured by timing a pressure step at its two ends. The leak of the closed gate is measured. The pressure drop of the shared supply line is measured against total flow. The outputs are the calibrated sensors, the volumes, the conductance as a function of setting with its uncertainty u_g, the thresholds, the delay and the line impedance.

37.6 Stage 1: Each Chamber Alone

With the gate closed, each chamber is run free for no fewer than 300 cycles after discarding the first 20. The period, the charging and discharging durations and the cycle-to-cycle jitter are recorded, together with the drift of the period over the session. The ensemble-average waveform $p(\theta)$, binned by phase, is fitted by weighted least squares to the piecewise-exponential model of Chapter 2, giving the effective supply pressure, the time constants and the open-valve target, and the structure of the residual is reported as a check on the first-order model. The fitted parameters give a predicted period by Equation (2.1), to be compared with the measured one.

The response curve is then measured directly. A pressure pulse of known volume is injected at 24 phases spread uniformly over the cycle, each phase repeated at least 10 times in random order, and the phase shift after at least three cycles of settling is read from the event timestamps. The response $\hat{Z}(\theta_j)$ is the shift per unit pressure perturbation. It is compared with the analytic $Z(\theta)$ computed from the fitted waveform parameters. The response data are not used in the fit, so this comparison is a prediction test of the phase response derivation, independent of coupling.

37.7 Stage 2: The Uncoupled Pair

With the gate closed, both chambers run together for no fewer than 300 cycles. The natural rate difference is estimated as $\hat{\delta} = T_1/T_2 - 1$ with its uncertainty from the period statistics. The shared-supply coupling is measured by comparing each chamber's period when the other runs and when it is held, and the leak coupling of the closed gate is measured from the modulation of the phase drift. The outputs are $\hat{\delta}$, a bound on residual coupling, and the supply cross effect. If $|\hat{\delta}|$ is larger than half the predicted locking range at the weakest gate setting, the supply of one chamber is trimmed to reduce it.

37.8 Stage 3: Coupled Runs

The gate is set to at least three conductance levels, chosen so that ϵT takes values near 0.02, 0.04 and 0.08. For each level the second chamber is started with a programmed delay so that the initial phase difference takes the seven values $0.05T$, $0.2T$, $0.35T$, $0.5T$, $0.65T$, $0.8T$ and $0.95T$, the actual value being read from the first events. Each condition is repeated at least five times in randomized order, blocked by session. Each run lasts the larger of six relaxation times $1/(2\epsilon)$ and $60T$. A single free-running chamber is run between blocks as a drift control.

To sample G at every phase, and not only near the equilibria, detuned runs are added at each level, with the supply of one chamber trimmed to produce $|\delta_d|$ at least 1.2 times the predicted peak of ϵG, so that the phase difference drifts through the whole cycle. Both signs of the detuning are run.

37.9 Estimating the Interaction Function

Three quantities must be kept apart. The *event lag* is the measured time $\lambda_k = t_{1,k} - t_{2,k}$ between matched valve-closing events of chambers 1 and 2 in the kth cycle, where the pairing

is by cycle count and not by nearest event. The *normalized lag* is λ_k divided by the period of the reference chamber measured over the same window, $t_{2,k+1}-t_{2,k}$ averaged over it; dividing by a nominal period instead biases the result by the relative frequency difference times the offset, which was 0.001 to 0.004 of a cycle in the simulations of this Part. The *model phase difference* is $\varphi=\theta_2-\theta_1$ of Chapter 30, in time units, an instantaneous quantity defined on the whole cycle.

The three are related but are not the same. At the closing of chamber 2 one has $\theta_2=T_2\equiv 0$, and chamber 1 reaches its own closing a time λ_k later, so the model phase difference at that instant is $\varphi(t_{2,k})=\lambda_k\big(1+O(\epsilon)\big)$, the correction coming from the rate of θ_1 over the lag, which is $1+Z_1 q_1$ and not 1. The averaged model of Chapter 29 describes the cycle-averaged phase difference, and the instantaneous phase at the closing differs from it by the oscillatory displacement W_1-W_2 of Chapter 34, also of order ϵ. The event-lag estimate of $\hat G$ therefore estimates the averaged interaction function only up to terms of first order in ϵ relative to the quantity measured, and the comparator simulation is asked to produce the same observable, the event lag from valve closings with the same pairing and normalization, as well as the model phase from the pressures, so that the two are compared with each other and not with the model phase alone.

The lag is *unwrapped* before any fit: the series λ_k is made continuous by adding integer multiples of the period so that successive values differ by less than half a period, since reduction modulo T puts a jump into the series whenever the drift crosses the origin of the cycle and a regression across that jump gives a spurious rate. A window of m consecutive cycles then gives a rate $\dot\lambda$ by linear regression of the unwrapped λ_k on time, and the rate of the model phase difference is taken to be $\dot\varphi=\dot\lambda$ at the order of the conversion above. The window is chosen with $m\geq 5$ and short enough that φ changes by no more than $0.05T$ across it. The estimate at the window's mean phase, read from the unwrapped lag, is

$$\hat G(\bar\varphi)=\frac{\dot\varphi-\hat\delta}{\hat\epsilon},\qquad \hat\epsilon=\hat g/\hat C. \tag{37.2}$$

Uncertainty. The phase marker is the valve-closing event, which occurs on the steep discharge branch. The timestamp variance of an event is

$$\sigma_t^2=\Big(\frac{\sigma_p}{|\dot p_l|}\Big)^2+\sigma_{\mathrm{valve}}^2+\frac{\Delta t_s^2}{12}\ \text{(omitted with interpolation)},\qquad |\dot p_l|=\frac{p_l-p_o}{\tau_o}, \tag{37.3}$$

where σ_p is the pressure noise, σ_{valve} the switching jitter and Δt_s the sample interval. The variance of one phase-difference value is $\sigma_\varphi^2=\sigma_{t,1}^2+\sigma_{t,2}^2+\sigma_{\mathrm{sync}}^2$, with σ_{sync} the clock offset between channels. For events at spacing T, the regression slope over m cycles has variance

$$\sigma_{\dot\varphi}^2=\frac{12\,\sigma_\varphi^2}{T^2\,m(m^2-1)}, \tag{37.4}$$

and the interaction function inherits

$$\sigma_{\hat G}^2=\frac{\sigma_{\dot\varphi}^2+\sigma_{\hat\delta}^2}{\hat\epsilon^{\,2}}+\hat G^2\Big(\frac{\sigma_{\hat\epsilon}}{\hat\epsilon}\Big)^2,\qquad \Big(\frac{\sigma_{\hat\epsilon}}{\hat\epsilon}\Big)^2=\Big(\frac{u_g}{g}\Big)^2+\Big(\frac{u_C}{C}\Big)^2. \tag{37.5}$$

The binned estimates are also fitted by a Fourier series to the fifth harmonic by weighted least squares, and confidence intervals are obtained by bootstrap over runs. Window length

trades variance against bias, and the bias from phase change across the window is checked by halving the window. The slope statistic is the coefficient a_1 of the fit $\hat{G} \approx a_1\varphi + b\,\varphi|\varphi| + a_3\varphi^3$ over $|\varphi| \le 0.1T$, with b and a_3 free, reported together with the same fit over $0.05T$ and $0.15T$ and with b fixed, for the reasons given in Section 37.3; its uncertainty is the bootstrap interval over runs.

37.10 Fitting and Prediction Kept Apart

The chamber, gate and detuning parameters are estimated from Stages 0 to 2 only. Before any Stage 3 run, the fitted parameters are frozen and the predictions $G_{\mathrm{pred}}(\varphi)$, the locking range, the stable offsets and the relaxation rates are written to a dated file. The primary test has no free parameter. A secondary test allows one free scale factor s on ϵ, fitted on a declared tuning half of the Stage 3 conditions and evaluated on the remaining test half; a value of s different from one, with its uncertainty, measures coupling efficiency that the model does not contain. Data used in a fit are never used to test the same fit.

37.11 Linearity and Averaging Checks

The derivation is first order in ϵ, so $\hat{G}$ should not depend on the gate level. The estimates at the three levels are plotted together against φ and tested for collapse. A systematic dependence on ϵ is fitted as $\hat{G}(\varphi;\epsilon) = G_0(\varphi) + \epsilon G_1(\varphi)$, and the size of G_1 against G_0 measures the dependence of the estimate on the gate level. That dependence has several possible causes: the error of first-order averaging, which the monograph has not been able to bound; gate nonlinearity, which Stage 0 measured; changes of the chamber waveform, thresholds or delay with the gate level; bias in the conversion from event lag to model phase described above; and calibration error in $\hat{g}/\hat{C}$. The observation alone does not isolate any one of them. The protocol separates them in two steps. The dependence reproduced by the comparator simulation of the fitted model at the same levels, which contains the averaging error and none of the unmodelled physics, is attributed to the model's own averaging error; the dependence of the bench estimate that the simulation does not reproduce is attributed to unmodelled effects, to be examined with the Stage 0 measurements, and is not called averaging error.

37.12 The Comparator Simulation and Its Convergence

The predictions are compared with a hybrid simulation of the fitted model, and that simulation has to be shown converged, and shown to agree with an independent solver, before its output is used as a comparator. Two solvers were used. The first is forward Euler with the valves switching at the first step that crosses a threshold, valve-closing times located by linear interpolation of the pressure across the step, and the delay held in a ring buffer with linear interpolation between samples. Interpolating the recorded closing times improves the timestamps but not the trajectory, since the dynamics before the switch are still used to the end of the step. The second is an event-aware fourth-order Runge–Kutta solver: each step is tested for threshold crossings, the crossing time is located by bisection on the step itself, the step is shortened to end at the crossing, the mode is switched there, and the integration

restarts from the switch, with the delayed pressure read from a cubic Hermite interpolant of the stored trajectory. The two share the model equations and the preparation, and nothing else. A step is accepted when halving it changes every compared quantity by less than a quarter of its measurement uncertainty.

Cycle A, Euler. The tested steps were 4×10^{-4}, 2×10^{-4}, 10^{-4}, 5×10^{-5}, 2.5×10^{-5} and 1.25×10^{-5}, that is $\Delta t/T$ from 4.4×10^{-4} to 1.4×10^{-5}. The first columns are the decay rate of a pair of diffusive-gate chambers at $\epsilon=0.08$ started at $0.06T$ (the tangent value 2ϵ is 0.160, and the first-order model fitted over the same window gives 0.1539), the antiphase offset of a delayed-gate pair at $\epsilon=0.08$, $\tau=0.512T$ with a supply difference of 0.1 (a rate difference of 0.0117), and the locked state of the loaded three-chamber network of Chapter 32.

Step	Pair decay rate	Delayed pair u/T	Network u_F/T	Network u_A/T	Network Ω
4×10^{-4}	0.085	0.6063	0.5079	0.0149	0.9969
2×10^{-4}	0.165	0.5987	0.5056	0.0153	0.9973
10^{-4}	0.154	0.6012	0.5060	0.0154	0.9974
5×10^{-5}	0.155	0.5996	0.5060	0.0158	0.9974
2.5×10^{-5}	0.155	0.5996	0.5065	0.0157	0.9974
1.25×10^{-5}	0.154	0.5998	0.5064	0.0158	0.9974

The coarsest step fails even with interpolation, because the valve switches up to one step late and that error is comparable to the phase differences being fitted. The step 2×10^{-4} gives 0.165 by chance, which shows that agreement at a single step is not evidence of convergence. From 10^{-4} downward the decay rate is 0.154 to 0.155, the delayed-pair offset is 0.5996 to 0.6012 with the period 0.90235 to 0.90224, and the network offsets vary by $0.0005T$ in u_F and $0.0004T$ in u_A. Taking the difference of the two finest steps as the numerical uncertainty, the Euler values are 0.154, 0.5997 ± 0.0002, $u_F/T=0.5064\pm0.0001$, $u_A/T=0.0158\pm0.0001$ and $\Omega=0.9974$.

Cycle A, event-aware solver. With the steps 4×10^{-3}, 2×10^{-3} and 10^{-3}, and in a separate test down to 5×10^{-4}, the locked network gave $u_F/T=0.5064$, $u_A/T=0.0158$ and $\Omega=0.99745$ at every step, and the uncoupled period reproduced the analytic value to 10^{-10}. These agree with the Euler values to $0.0001T$ and 0.0001. The offsets under extra detuning of F by $+0.010$, $+0.016$ and -0.010 were 0.5812, 0.6487 and 0.4269. The drum-tracking errors agreed with those of the Euler solver to 0.001 radian. A scan of 61 release ends gave the same final states and classes as the Euler scan, with settling times agreeing to 0.06 time unit. The event-aware solver has no timing noise at the level of 10^{-4}, so it can follow a small phase difference at weak coupling, which Euler at the steps used could not. The decay rate of a pair of diffusive-gate chambers divided by ϵ was 1.926 to 1.934 for a start at $0.06T$, 1.976 for a start at $0.02T$ and 1.994 for a start at $0.005T$, and the same to within 0.01 for ϵT from 0.01 to 0.16. The ratio therefore depends on the amplitude of the displacement, as the nonlinearity of G requires, and not on the coupling, and it tends to 2 as the amplitude falls.

Limits of locking. Fixed-time bisection on the extra detuning, with runs of 400 to 500 time units, is not reliable near the fold, because beyond the fold the phase slips only after a

long passage through the bottleneck: for a delayed pair in cycle A it gave $[-0.0140, +0.0130]$ with one solver and run length and $[-0.0130, +0.0145]$ with another, while a run of 600 time units at a detuning of $+0.0142$, which the second bisection had classified as locked, was clearly slipping. The limits below were therefore bracketed with runs of 3000 time units, between the largest detuning that stayed locked in the last 800 time units and the smallest that slipped, and in cycle A they were checked by ramping the detuning slowly and extrapolating the loss point to zero ramp rate as $r^{2/3}$, which gave $[-0.0152, +0.0199]$ for the network and $[-0.0127, +0.0130]$ for the pair, within 0.0001 of the brackets. The ramp extrapolation was erratic in cycle B at weak coupling and is not used there. In every row the brackets coincide to 0.0001. A bracket certifies lock inside it for the initial condition used and does not rule out a locked state outside the basin of that condition. The hybrid limits are compared with the averaged-model limits for the network and, for the pair, with the first-order limits computed from the pair interaction tables with the delay phase taken at the collective rate, the pair counterpart of the averaged network model, with the coupling varied by a common factor.

Cycle	ϵT_{SF}	Network, hybrid	Network, averaged	Pair, hybrid	Pair, first order
A	0.027	$-0.0079, +0.0095$	$-0.0080, +0.0096$	∓ 0.0064	∓ 0.0062
A	0.054	$-0.0151, +0.0198$	$-0.0156, +0.0199$	$-0.0126, +0.0129$	$-0.0124, +0.0128$
A	0.109	$-0.0286, +0.0406$	$-0.0305, +0.0403$	$-0.0246, +0.0258$	$-0.0248, +0.0258$
A	0.163	$-0.0409, +0.0616$	$-0.0450, +0.0609$	$-0.0359, +0.0387$	$-0.0364, +0.0388$
B	0.082	$-0.0085, +0.0076$	$-0.0094, +0.0100$	$-0.0074, +0.0078$	$-0.0076, +0.0080$
B	0.109	$-0.0105, +0.0124$	$-0.0121, +0.0156$	$-0.0095, +0.0102$	$-0.0100, +0.0106$
B	0.163	$-0.0140, +0.0207$	$-0.0174, +0.0258$	$-0.0132, +0.0149$	$-0.0144, +0.0158$
B	0.217	$-0.0167, +0.0278$	$-0.0224, +0.0354$	$-0.0163, +0.0192$	$-0.0184, +0.0212$
B, no loads	0.082	$-0.0078, +0.0136$	$-0.0087, +0.0141$	$-0.0074, +0.0078$	$-0.0076, +0.0080$
B, no loads	0.163	$-0.0136, +0.0262$	$-0.0169, +0.0281$	$-0.0132, +0.0149$	$-0.0144, +0.0158$

Every percentage in this paragraph is the difference between a first-order prediction and the hybrid limit divided by the predicted limit, as defined at the start of the diagnosis in Chapter 34; the prediction is the frequency-only first-order one (P1) unless stated, and for the pair it is the first-order value from the pair tables. In cycle A the hybrid limits equal the averaged and first-order ones to within 0.0001 at the weakest coupling and depart from them progressively as the coupling grows, to 9 percent at the lower limit of the network at the strongest, which is the trend a first-order error would show, although the trend alone does not identify the cause; the pair is within 3 percent throughout. In cycle B the pair agrees with the first-order limits to 3 percent at the weakest coupling tested and departs by 5, 8 and 11 percent at the lower limit as the coupling grows to 0.217, again the trend a first-order error would show. The lower limit of the loaded network follows the same pattern, from 10 to 25 percent. The upper limit of the loaded network is 20 to 24 percent below the averaged one over the whole tested range with no trend, whereas without loads it is within 4 to 7 percent. Hybrid single-chamber load shifts agree with the first-order formula to 0.00003, so the discrepancy is not an error of the constant shifts used. The causes are examined in Chapter 34: freezing the bladder dynamics changes the limits by at most 0.0002, recomputing the interaction for the detuned chambers removes a few percent, and the second-order term as computed does not account for the remainder. The apparent absence of a trend with coupling arises because the load spread was held fixed while the coupling was reduced; with the spread scaled with the coupling the discrepancy falls roughly in proportion to it, which is consistent with an averaging error but does not identify the cause.

Cycle B. The Euler steps tested were 6×10^{-4}, 3×10^{-4}, 1.5×10^{-4}, 7.5×10^{-5}, 3.75×10^{-5} and 1.875×10^{-5}, the same $\Delta t/T$ as for A to within a factor of 1.3. The loaded locked state of the retuned network had u_F/T from 0.5097 at the coarsest step to 0.5094 at the four finest, u_A/T between 0.0131 and 0.0134, equal to 0.0131 at the finest, and Ω from 0.9958 to 0.9960. The event-aware solver gave 0.5094, 0.0131 and 0.99604 at its two steps. Its simulated offsets under extra detuning of F at $+0.010$, $+0.020$ and -0.010 were 0.5374, 0.5894 and 0.4800. The fitted relaxation rate of the displaced free leg was 0.256 ± 0.004. A scan of 92 release ends gave the same final states in both solvers.

Slow states and run length. For the two unretuned settings that have no state near antiphase, a run of 260 time units shows a transient and not the state: the state with the original physical delay read $u_F/T = 0.025$ and $u_A/T = 0.179$ at that time and 0.221, 0.347 after 3000 time units, and the window-tuned delay read 0.554 and 0.178 and then 0.663, 0.276. These attract at rates of order 0.015, so reading them requires runs of several hundred time units beyond the settling time; they are close to the stable states of the averaged model. Euler at the steps 6×10^{-4} to 1.875×10^{-5} also gave different values at 260 time units, for the same reason. The half-period setting is settled at 260 time units and read $u_F/T = 0.535$, $u_A/T = 0.048$.

Offset convention, event timing and the bladder start. Three preparation choices change the answer by as much as the model differences being assessed. The first is the convention for the offset. The averaged model computes phase offsets $x_i - x_S$ in cycles of the chambers, and a hybrid offset must be the same quantity. The lag between valve closings divided by the *nominal* period is not: chambers with different natural periods, here by the load shifts, then differ by the product of the relative frequency difference and the offset. That convention gave $u_F/T = 0.5051$ and $u_A/T = 0.0132$ in cycle A and 0.5074, 0.0092 in cycle B, against 0.5064, 0.0158 and 0.5094, 0.0131 with the lag divided by the period of the reference chamber, which agrees with the phase read from the pressures to 10^{-4}. The identification used throughout is therefore explicit: the lag divided by the locked period of the reference chamber is a cycle fraction, and it is identified with the model phase offset $x_i - x_S$ read at the closing of S, which is the instantaneous phase and differs from the cycle-averaged offset $y_i - y_S$ by $W_i - W_S$ (Chapter 34); the identification with the pressure-read phase holds to 10^{-4} in the simulations and is a convention to be checked, not an identity. The difference, $0.0013T$ to $0.0039T$, is as large as the discrepancy between the averaged and hybrid offsets that the first convention suggested, and accounts for all of it in cycle A, where with the correct convention the averaged and hybrid offsets agree to $0.0002T$. The second is event timing: a simulation that takes event times at step boundaries is biased by the step. The third is the bladder start. With the bladders started at pressure one and run for 150 time units at the step 10^{-4}, cycle A gave $u_F/T = 0.5054$, $u_A/T = 0.0145$ and $\Omega = 0.9972$, and after 1500 time units 0.5060, 0.0154 and 0.9974, which equals the value from a run started at the periodic mean. Bladder time constants are 33 to 133 time units in cycle A and 65 to 260 in cycle B, so runs started away from the periodic state are not settled in the window used to read offsets. The comparator must therefore start every slow state at its periodic mean and run long compared with its slowest time constant, and the protocol requires the same of the bench: bladders must be pressurized to their working mean before a run is read. The same

first-order load shift formula needs the bladder at its periodic steady state, and evaluated with a shorter spin-up it overstates small loads (Chapter 30).

Release and recapture are more solver-sensitive in cycle B than in cycle A, but only in the route taken. The final states and the classification of every scanned release agree between the solvers, and the settling times agree to 0.06 time unit in A and to 0.03 at the median in B. In B one release, with recapture phase $0.94T$, near the boundary where the route around the cycle reverses, went the opposite way round in the two solvers and settled in 113 and 67 time units, so the location of that boundary is solver-dependent at the level of $0.001T$ in the recapture phase (Chapter 33).

What was not checked. Both solvers share the model: ideal threshold switching with no valve dynamics, linear gates and bladders, an ideal gas treated as a single capacitance, and a pure transport delay. Agreement between them tests the numerical solution of the model and says nothing about the model. The integration of the drum commands and the sampling of the pressures for the drum analysis, at intervals of 2×10^{-3} for cycle A and 4×10^{-3} for cycle B, were not varied. The cause of the cycle B limit discrepancy is only partly established (Chapter 34).

37.13 Reporting and Decision Rules

The principal report is a table with, for each phase bin, $\hat{G}$, its standard uncertainty, G_{pred}, the difference and the difference in units of the combined uncertainty. The discrepancy statistic is $D = \mathrm{RMS}(\hat{G} - G_{\mathrm{pred}}) / \max |G_{\mathrm{pred}}|$ with its bootstrap interval. The following agreement criteria are proposed and are to be fixed before the runs and not after: the slope coefficient a_1 equals the value -2, or the value from the same fit to the comparator simulation, within two combined standard uncertainties; at least 90 percent of bins lie within two combined standard uncertainties of the prediction; and $D \le 0.15$. The thresholds are design choices and carry no claim.

Disagreement is interpreted by its form. A shape error at every gate level points to an error of the waveform or the response curve, and Stage 1 locates it. An error that grows with ϵ points to averaging error or gate nonlinearity. An even component in $\hat{G}$, which the symmetric prediction does not have, points to a chamber mismatch or an unmodeled one-way effect. A constant offset of $\dot{\varphi}$ across phases points to a residual detuning. A failure of the slope test is the strongest negative result for the derivation, but it is a failure of the derivation, the linearity of the gate and the calibrations of g and C together; it is assigned to one of them only by the Stage 0 measurements.

37.14 Later Stages and Limits

The same stages are to be repeated for the delayed gate, with the delay calibrated and the window of Chapter 30 tested, for the one-way gate, and finally for three chambers with a drum interface, each with the added calibrations it requires. The protocol does not test loads on the chambers, the shared-supply coupling beyond the cross-effect measurement of Stage 2, joint mechanics or balance, and a pass would support the interaction-function derivation for the tested cycle and gate only.

Chapter 38

Unresolved Design Choices

The architecture is specified at the level of parts and roles, and the following choices remain open. They are collected here so that the conceptual construction is not mistaken for a tested design.

Item	What is unspecified
Shell	Form (folded, laminated, molded), thickness, seam construction, allowable stress
Reinforcement	Rib layout, cell geometry, wall strength against sealing requirements
Guides	Mounting method, eye material, spacing, routing at each joint
Attachments	Patch area, load spreading, method of termination on the shell
Blackbox	Internal mechanism, spool or drop, number of channels, location (concentrated or segmental)
Pneumatic network	Chamber geometry, seals, gate construction, operating pressures, source type
Pressure role	Which of stiffening, return, take-up drive or cycling oscillator is used where
Compressed-air reservoirs	Capacity, release trigger, discharge location
Fabrication route	Which of the four forming concepts, in what combination, and in what order
Assembly order	Sequence of forming, cord installation, sealing and testing

Of these, the pressure role and the guide mounting have the largest effect on the theory, since the first determines how pressure enters the torque relation and the second determines the moment arms and friction losses of the cord layer.

Chapter 39

Conclusion

This monograph proposed that coordination in a lightweight robot can be treated as the establishment, selective release and re-establishment of phase relationships among oscillators arranged in a hierarchy, and it worked out what that claim does and does not support. The mathematical results are proved within a stated phase model. In that model a group of coupled oscillators keeps its internal coordination within an explicit bound set by the coupling margin, the frequency dispersion and the external coupling, for every schedule of the external couplings and, with gated adaptation, for every adaptation of the internal weights that stays inside a certified box. Collective phases obey an effective-oscillator equation with a bounded remainder, so that a group can be released and recaptured: for two groups with uniform or nonuniform coupling, and for a three-level hierarchy with uniform tiers, the retained groups keep their coordination, the released group advances at its own rate up to residual paths, and recapture brings the relative phase into an explicit arc, finite in time for any enlarged arc and asymptotic for the exact one. For three groups the internal coordination is preserved under any gate schedule, release and recapture follow pairwise when only one link is active, and simultaneous recapture holds under consistent targets and a margin condition. These results are sufficient conditions, stated with their excluded region, and some of them are conservative, as the simulations showed.

The physical realization is conditional. A first embodiment keeps the tendon and pressure channels independent, places the oscillators in the pneumatic layer, and fixes the couplings by gate conductances whose first-order effect was derived from the chamber dynamics. That derivation gave the interaction functions of a diffusive gate, a half-period delayed gate and a one-way gate, with their locking ranges, stability and sensitivity to load, and showed that the sinusoidal approximation of the oscillator theory is accurate for some cycles and quantities and not for others. Residual paths through leakage, the shared supply and the body were modeled as bounded perturbations, with the caveat that the shared supply is really a state-dependent coupling. A phase-to-cord-length interface with an explicit error budget carries the pneumatic coordination to the drum coordinates.

The simulation evidence is limited and is described as such. A three-chamber gait network with stance and arms in phase and the free leg in antiphase was solved for its locked state, shown stable by linearization, and checked against hybrid simulations with loads, detuning, a shared supply, release and recapture, and drum tracking, in two pneumatic cycles. The behavior survived a change of cycle only after recalibration: the physical delay failed, the

half-period delay worked, and the couplings had to be raised to restore the margin. The simulations were checked for dependence on the step, on the initial state of the bladders, on the solver and on the convention for the offset. With offsets measured as phases, the first-order averaged model reproduces the locked offsets of the hybrid simulation to $0.0002T$ in the first cycle and $0.0012T$ in the second, and its locking limits to a few percent at weak coupling, with departures that grow with the coupling. In the second cycle the predicted limits of the loaded network are about twenty percent too wide at the retuned coupling for the frequency-only first-order prediction, and fourteen to seventeen percent too wide when the interaction is recomputed for the detuned chambers (percentages of the predicted limit), a discrepancy for which the corrections tested (the bladder dynamics, the waveform change of the detuned chamber, the offsets and the second-order averaging term as computed) are insufficient, and whose cause is unresolved. Agreement of the locked offsets, which are interior solutions, does not establish agreement of the boundaries of locking, and the two are reported separately. Within the stated model, the unaveraged phase equation of a single-pressure chamber is exact under Assumption W, and that of a loaded chamber is exact under transversality with a second, slow variable whose measured effect is small, so for a single-pressure chamber the distance between the phase model and the chamber is the averaging error, and for a loaded chamber it also includes the neglected slow variable. The finite-horizon averaging bound proved is first order in the coupling at a fixed time and has constants too large to be informative at the couplings used, its persistence statement is conditional on hypotheses that fail at those couplings, and its extension to loaded chambers is open. Recapture in the second cycle returned to the intended state from every tested initial condition but was slow from phases near in-phase. The averaged predictions are first order in the coupling; the hybrid simulations integrate the stated chamber equations and are not first order. All are calculations within the model and not measurements of a body, and the validation program of Part V has not been run. Hardware precedents for coupled rhythm generators exist and are cited in the preface, and one of them is a reminder that components communicating is not the same as a robot behaving well, which is why the evidence here is stated at the level of the phase model and the drum coordinate and no further.

What the monograph leaves open is substantial. The passage from drum coordinates to joint motion depends on routing, slack, elasticity and load, and nothing here establishes joint angles; balance and support require mechanical and contact conditions that the oscillator theory does not supply, and the stance-admissibility hypothesis stays a requirement on the body. The frustrated case, in which the preferred offsets do not close around a cycle, is not treated, hierarchies deeper than three levels and nonuniform tiers at three levels are not proved, adaptation of the cross couplings is open, and the geometric reading of the summary as a motion-generating amplitwist is unproved. The design choices of Part V, from shell thickness to valve construction, are unspecified. The central claim, that control consists partly in changing which movements constrain one another, at which level and for how long, is supported within the model and is a hypothesis about the body.

Appendix A

Proposed Biological Extensions

The earlier proposals for the endomarionette include a biological-materials branch: a yogurt-based interstitium, internally grown SCOBY material, leather-like reinforcement, and ligaments and tendons knitted internally as yarns. These were proposed as extensions of the paper body. They have not been built or experimentally demonstrated, and nothing in the main text depends on them. They are recorded here so that the options remain visible.

Appendix B

Ledger of Proof Obligations

Each obligation is assigned a category and a status. The categories are those of the scope chapter: identity, theorem, approximation, numerical finding and proposal, with open for what is not established. The ledger describes the draft as of its date and should be revised as proofs are completed and tests are run. The three most consequential obligations are a sharp bound on the averaging error of the chamber phase equations, the physical realizability of the coupling law, and the passage from drum coordinates to joint mechanics and balance.

Obligation	Category	Status
Unaveraged phase equation of a single-pressure chamber	Theorem	Exact under Assumption W (Theorem 34.1); the averaged interaction equation is an approximation to it. Checked numerically to 10^{-11} in the median.
Finite-horizon averaging bound and persistence of a locked state	Theorem (repaired), conditional proposition	The finite-horizon bound (Theorem 34.2) is first order in ϵ' at a fixed time, with an order-one contribution from the unfinished window, and of order ϵ'^2 only at fixed integer times, not over the averaging horizon $1/\epsilon'$; an earlier version claimed second order at every time and was wrong. Proposition 34.1 is conditional: its contraction neighborhood is not certified and its hypothesis $\Delta \leq \rho_0/4$ fails at the couplings used (the formula evaluates to about 15 and 145 cycles for $K_c = 1$, above any admissible ρ_0). The constants are uninformative by four to five orders of magnitude against observed offset differences under 0.003.
Sharp bound on the averaging error of the chambers	Open	Highest mathematical priority. Needs a second-order window-map bound with remainder; the lower limit of locking and the larger discrepancy of the loaded network are unexplained.

Obligation	Category	Status
Second-order averaging estimate	Numerical finding	Evaluated numerically; does not account for the cycle B limit discrepancy; not proved for the discontinuous right-hand side.
Exact equations of a loaded chamber	Theorem	Two variables per chamber, with a reset (Theorem 34.3), under transversality. The phase-only model obtained by freezing the slow variable at its fixed point is an approximation, not an invariant subsystem. Slow-variable sensitivities were measured at one fixed point, and the estimate of the slow variable's deviation is a local linear estimate, not a bound. Extension of the averaging bound to this form is open.
Shift lemma, coherence adequacy, two-sided relation between coherence and spread	Identity, lemma	Proved.
Two-oscillator locking and group cohesion, uniform and through the interconnection graph	Theorem	Proved as sufficient conditions, invariant and attracting parts stated separately. Standard in type; not a new method.
Effective oscillator, one level, with explicit remainder	Theorem	Proved for uniform and nonuniform coupling.
Selective release and recapture of two groups	Theorem	Proved. Capture is finite-time into any enlarged arc and asymptotic for the exact arc, with a uniform time bound only away from the excluded band, which contains the unstable equilibrium.
Three-level hierarchy with uniform tiers	Theorem	Proved. The estimate degrades geometrically with depth; this is a property of the estimate and not shown to be intrinsic. Nonuniform tiers and deeper hierarchies are open.
Adaptation of internal weights	Theorem	Proved for gated adaptation inside a certified box, at any rate. Adaptation of cross weights and a fatigue law that lowers the floor are open.
Moving targets and drifting natural frequencies	Corollary	One-line consequence of the theorems.
Three groups: internal preservation, pairwise-exclusive schedules, simultaneous recapture	Theorem	Proved under consistent targets and a margin condition. The frustrated case is open and deferred.
Directed cascade between segments	Theorem	Local exponential stability proved; global behavior, nonuniform coupling and the pneumatic one-way chain are open.
Sensory correction	Corollary	Sufficient condition from an added bounded disturbance; conservative.
Transition mechanisms	Proposal	Four hypotheses with discriminating observations; the locking-preserving case follows from capture.

Obligation	Category	Status
Stopping and restarting; amplitude and offset commands	Theorem, proposal	Sufficient condition for the engagement model and an explicit transient for critically damped amplitude commands; mechanical re-engagement open.
Hybrid consistency across resets	Proposition, open	Conditional invariance across resets, with a restrictive worst-case bound on spread widening; the form of a foot-contact reset is open.
Amplitwist interpretation of summaries; compatibility of composition with reduction	Open	Interpretation only. The essay's definition is recorded and the identification with a motion-generating map is unproved.
Limits of coherence and phase summaries	Example	Example given, theorem open.
Pneumatic relaxation cycle: existence, attraction, dwell time	Theorem	Proved for the idealized piecewise-linear model.
Interaction functions of two chambers for a diffusive, a delayed and a one-way gate	Approximation	First order in the coupling. The slope result, the symmetry statement and the parity rule are proved in the model; locking ranges and offsets are numerical and were checked against hybrid simulation at two cycles.
Accuracy of the sinusoidal approximation to the interaction function	Numerical	Cycle dependent: adequate for the locking range in one cycle, poor in the other. A general statement is open.
Residual coupling through leakage, shared supply and body	Theorem, open	Bounded-perturbation model proved for two groups. The shared supply acts as a state-dependent coupling, and a uniform bound that would make the representation valid has not been established.
Interface from chamber phase to cord length	Theorem	Error budget proved for a first-order drum in steady sinusoidal operation. The map from drum position to joint motion is open.
Worked three-chamber gait network	Numerical	Linear stability by linearization, margins and corner checks numerical, hybrid simulations at one parameter set per cycle; no theorem. Corner checks stay numerical unless an argument shows the corners cover the detuning box.
Second cycle and recalibration	Numerical	The physical delay failed, the half-period delay worked, and the couplings had to be raised. No rule for recalibration is claimed.

Obligation	Category	Status
Error of first-order averaging at large coupling, and treatment of the delay	Open	The delay is treated as a quasi-static shift, which the finite-horizon bound covers through the factor $1 + \bar{d}$ (Theorem 34.2). At the coupling of the second cycle the frequency-only first-order limits of locking are about twenty percent too wide for the loaded network, as a percentage of the predicted limit, and fourteen to seventeen percent with the interaction recomputed (Chapter 34).
Cord-count bound for tendon control	Theorem	Known result, short proof included.
Shell stability, reachable joint motions, mechanical feasibility	Open	Proposal level only. Design choices are tabulated in Part V.
Stance admissibility and balance	Open	Stated as a hypothesis on the body mechanics.
Locomotion by pressure alone; redistribution by timed releases	Open	Untested hypotheses.
Robustness to leaks, friction, delay, shell deformation and noise	Open	Not bounded.
Biological extensions of the body	Proposal	Recorded in an appendix; nothing depends on them.
Validation program	Proposal	Specified, not run.
Bench protocol for two chambers	Proposal	Executable protocol (Chapter 37), not run; its output is the measured interaction function and its disagreement with the prediction.
Step-size and solver convergence of the hybrid simulations	Numerical finding	Euler at six steps and an independent event-aware fourth-order solver agree on offsets to $0.0003T$ and on settled final states in all scanned releases; locking limits bracketed to 0.0001 (Chapter 37). Both solvers share the model, so agreement says nothing about the model.
Slope prediction $G'(0) = -2$	Identity	Holds in the first-order averaged equation under the scalings of Section 37.3; its bench test also depends on calibrations of g and C and on gate linearity.

Obligation	Category	Status
Averaged model against converged hybrid simulation	Numerical finding	Locked offsets agree to $0.0002T$ in cycle A and $0.0012T$ in cycle B once offsets are measured as phases. Locking limits agree within a few percent at weak coupling and depart as coupling grows; in B the hybrid limits of the loaded network are 19 to 20 percent inside the frequency-only prediction and 14 to 17 percent inside the recomputed one at the retuned coupling (percent of the predicted limit); the tested corrections are insufficient to account for this and the cause is unresolved. Agreement of locked offsets does not establish agreement of the locking boundaries (Chapter 34). No error bound is proved.
Related mechanisms (origami, tendon, McKibben)	Proposal	Precedents only (Chapter 28), verified at the level of abstracts and records and not full text; McKibben muscles are an alternative, not part of the first embodiment.

The principal theorem is proved for two groups with uniform coupling (Theorem 12.1), for two groups with nonuniform coupling through the internal graph (Theorem 13.3), and for a three-level hierarchy with uniform tiers (Theorem 15.1). The cohesion theorems establish a bounded spread, in an invariant set and an attracting set, and not the existence of a phase-locked solution or convergence to one; with time-dependent drives or adapting weights a strictly locked solution need not exist, and none is claimed.

Bibliography

[1] J. Cheslet, R. Beaubois, T. Duenki, *et al.* Biomimetic Snake Locomotion Using Central Pattern Generators Network and Bio-Hybrid Robot Perspective. *Artificial Life and Robotics*, 29:479–485, 2024. doi:10.1007/s10015-024-00969-0. Only the first three authors are verified here, and the full text could not be retrieved, so the description in the text rests on the abstract and secondary records.

[2] F. Dörfler and F. Bullo. Synchronization in Complex Networks of Phase Oscillators: A Survey. *Automatica*, 50(6):1539–1564, 2014.

[3] J. Duysens and H. W. A. A. Van de Crommert. Neural Control of Locomotion; Part 1: The Central Pattern Generator from Cats to Humans. *Gait & Posture*, 7(2):131–141, 1998.

[4] M. Golubitsky, I. Stewart, P.-L. Buono, and J. J. Collins. Symmetry in Locomotor Central Pattern Generators and Animal Gaits. *Nature*, 401:693–695, 1999.

[5] S. Grillner. Biological Pattern Generation: The Cellular and Computational Logic of Networks in Motion. *Neuron*, 52(5):751–766, 2006.

[6] A. J. Ijspeert. Central Pattern Generators for Locomotion Control in Animals and Robots: A Review. *Neural Networks*, 21(4):642–653, 2008.

[7] Y. Kuramoto. *Chemical Oscillations, Waves, and Turbulence*. Springer, Berlin, 1984.

[8] C.-P. Chou and B. Hannaford. Measurement and Modeling of McKibben Pneumatic Artificial Muscles. *IEEE Transactions on Robotics and Automation*, 12(1):90–102, 1996.

[9] S. Felton, M. Tolley, E. Demaine, D. Rus, and R. Wood. A Method for Building Self-Folding Machines. *Science*, 345(6197):644–646, 2014.

[10] J.-J. Lee and L.-W. Tsai. The Structural Synthesis of Tendon-Driven Manipulators Having a Pseudotriangular Structure Matrix. *International Journal of Robotics Research*, 10(3):255–262, 1991.

[11] E. Marder and R. L. Calabrese. Principles of Rhythmic Motor Pattern Generation. *Physiological Reviews*, 76(3):687–717, 1996.

[12] R. V. Martinez, C. R. Fish, X. Chen, and G. M. Whitesides. Elastomeric Origami: Programmable Paper-Elastomer Composites as Pneumatic Actuators. *Advanced Functional Materials*, 22(7):1376–1384, 2012.

[13] H. Nakao. Phase Reduction Approach to Synchronisation of Nonlinear Oscillators. *Contemporary Physics*, 57(2):188–214, 2016.

[14] A. Pikovsky, M. Rosenblum, and J. Kurths. *Synchronization: A Universal Concept in Nonlinear Sciences*. Cambridge University Press, Cambridge, 2001.

[15] P. Rothemund, A. Ainla, L. Belding, D. J. Preston, S. Kurihara, Z. Suo, and G. M. Whitesides. A Soft, Bistable Valve for Autonomous Control of Soft Actuators. *Science Robotics*, 3(16):eaar7986, 2018.

[16] D. Rus and M. T. Tolley. Design, Fabrication and Control of Origami Robots. *Nature Reviews Materials*, 3:101–112, 2018.

[17] H. Sakaguchi and Y. Kuramoto. A Soluble Active Rotator Model Showing Phase Transitions via Mutual Entrainment. *Progress of Theoretical Physics*, 76(3):576–581, 1986.

[18] S. Shirasaka, W. Kurebayashi, and H. Nakao. Phase Reduction Theory for Hybrid Nonlinear Oscillators. *Physical Review E*, 95:012212, 2017.

[19] S. H. Strogatz. From Kuramoto to Crawford: Exploring the Onset of Synchronization in Populations of Coupled Oscillators. *Physica D*, 143(1–4):1–20, 2000.

[20] S. H. Strogatz. *Nonlinear Dynamics and Chaos*, 2nd edition. Westview Press, Boulder, 2015.

[21] B. Van der Pol. On Relaxation-Oscillations. *Philosophical Magazine*, 2(11):978–992, 1926.

[22] W. Wang, J. Guo, Z. Wang, and G. Xie. Neural Controller for Swimming Modes and Gait Transition on an Ostraciiform Fish Robot. In *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, pages 1564–1569, Wollongong, 2013.

[23] M. Wehner, R. L. Truby, D. J. Fitzgerald, et al. An Integrated Design and Fabrication Strategy for Entirely Soft, Autonomous Robots. *Nature*, 536:451–455, 2016.

[24] A. T. Winfree. Biological Rhythms and the Behavior of Populations of Coupled Oscillators. *Journal of Theoretical Biology*, 16(1):15–42, 1967.

[25] Flyxion. Hierarchical Continuation and the Geometry of Locomotion. Unpublished essay, July 2026.

[26] Flyxion. From Classifier to Inverse Engine: Gesture Recognition as Latent Embodied Semantic Reconstruction. Unpublished essay, May 2026.

[27] Flyxion. Wireheading and the Conditions of a Common World. Unpublished essay, October 2026.