

Doctoral thesis

Natural Locomotion Families

Principle, Method, and Control

Mirado Mortel

Thesis director: Luc Jaulin

Co-supervisors: Lionel Lapierre and Simon Rohou

Résumé

Un locomoteur articulé convertit des mouvements internes en translation ou rotation par son interaction mécanique avec l'environnement. Plutôt que de prescrire une chorégraphie articulaire, cette thèse cherche les mouvements qu'organise sa dynamique passive, ou intrinsèque, puis comment les maintenir sans en effacer la structure. Le point de départ est l'*oscillation naturelle*. Un mode linéaire l'approxime près d'un équilibre, mais, à amplitude finie, fréquence, forme d'onde et coordination varient avec l'énergie; les modes non linéaires deviennent nécessaires.

La périodicité est une boucle fermée dans un espace d'état donné. Si cet état inclut la pose du corps, sa fermeture interdit tout déplacement net. Pour avancer, seul l'état mécanique réduit doit former la boucle : les déformations, leurs vitesses et les vitesses de translation et de rotation du segment principal—le porteur—reviennent exactement, tandis que sa pose accumule un déplacement net. Ce mouvement est relatif-périodique. La mécanique géométrique de la locomotion calcule le déplacement produit par une boucle interne prescrite, mais ne montre pas que la dynamique passive sélectionne cette boucle et son cadencement. Un cycle naturel est donc une solution non forcée avec un retour réduit exact et un incrément de pose différent de l'identité. Sous des hypothèses de rang et de plongement, restaurer la phase d'une branche régulière donne une carte locale du *Natural Locomotion Manifold* (NLM), sans postuler de variété globale.

La méthode étudie des locomoteurs plans à deux et trois segments sous contraintes idéales de vitesse, qui restreignent le mouvement sans fournir d'énergie. Les vitesses autorisées forment leur noyau; la projection des équations de Lagrange–d'Alembert y élimine les réactions sans travail. Éliminer temporairement les vitesses du porteur expose un oscillateur transverse effectif. Les réintroduire réouvre le canal propulsif, qui transporte le corps et échange énergie et puissance avec cet oscillateur; l'intégration de la pose reste distincte. Pour le locomoteur à deux segments, la contribution analytique centrale est un théorème d'équivalence démontré. Après le retour des variables du porteur, l'oscillateur effectif retrouve sa vitesse sur la section choisie si et seulement si son échange net d'énergie avec le canal propulsif est nul sur le cycle. Sous les hypothèses énoncées—vitesse non nulle, section transverse orientée et stockage injectif—cette équivalence fournit la condition scalaire de retour manquante, non l'existence, la stabilité ou la régularité d'une branche. Avec trois segments, un niveau d'énergie ne détermine plus l'état transverse vectoriel; le retour complet doit précéder la classification des familles issues des oscillations en phase et en opposition de phase. Une référence rectiligne appariée évite d'attribuer toute la translation reconstruite à l'oscillation interne.

Avec dissipation, une réponse maintenue doit revenir au même état et recevoir sur un cycle autant de travail que l'énergie dissipée. Une continuation ajoute pertes et actionnement autour d'un cycle; une étude distincte organise les réponses synchronisées dans l'espace orbite–forçage. Seules leurs composantes régulières et locomotrices donnent des cartes locales du *Resonant Locomotion Manifold* (RLM). Une boucle rapide régule localement la phase, tandis qu'une correction plus lente régule la vitesse moyenne. La thèse relie ainsi oscillation naturelle, locomotion conservatrice, maintien résonant et commande dans le cadre local énoncé.

Mots-clés. locomotion mécanique; oscillations naturelles; modes non linéaires; mécanique géométrique; mouvement relatif-périodique; contraintes non holonomes; échange d'énergie; réponse forcée; résonance; boucle à verrouillage de phase.

Abstract

An articulated locomotor converts internal motion into body translation or rotation through mechanical interaction with its environment. Rather than prescribing a joint choreography, this thesis seeks the motions organized by the mechanism's passive, or intrinsic, dynamics, then asks how to maintain them without erasing their structure. The starting point is a *natural oscillation*. A linear mode approximates it near equilibrium, but at finite amplitude natural frequency, waveform, and coordination vary with energy; nonlinear modes are then required.

A natural oscillation is not yet locomotion. Periodicity is a closed loop in a specified state space. If that state includes body pose, closure prevents any net displacement. To advance, only the reduced mechanical state must form the loop: internal deformation and rates, together with the forward and yaw velocities of the main segment—the carrier—return exactly, while its pose accumulates a net displacement. This motion is relative-periodic. The geometric mechanics of locomotion reconstructs the displacement produced by a prescribed internal loop but does not show that the passive dynamics selects that loop and its timing. A natural cycle is therefore an unforced solution with exact reduced return and a nonidentity pose increment. Under rank and embedding hypotheses, restoring the phase of a regular branch gives a local chart of the *Natural Locomotion Manifold* (NLM), without assuming a global manifold.

The method studies planar two- and three-segment locomotors under ideal velocity constraints, which restrict motion without supplying energy. The allowed velocities form their kernel; projecting the Lagrange–d'Alembert equations onto it removes the no-work reactions. Temporarily eliminating the carrier velocities exposes an effective transverse oscillator. Restoring them reopens the propulsive channel, which transports the body and exchanges energy and power with that oscillator; pose integration remains distinct. For the two-segment locomotor, the central analytical contribution is a proved equivalence theorem. Once the carrier variables have returned, the effective oscillator recovers its speed on the chosen section if and only if its net energy exchange with the propulsive channel over the cycle is zero. Under the stated hypotheses of nonzero speed, a transverse oriented section, and injective storage, this equivalence supplies the missing scalar return condition, not branch existence, stability, or regularity. With three segments, one scalar energy level no longer determines the vector transverse state; full return must precede classification of families continued from in-phase and anti-phase oscillations. A matched straight reference prevents the entire reconstructed translation from being attributed to internal oscillation.

With dissipation, a maintained response must return to the same state and receive over one cycle as much work as the losses dissipate. One continuation adds losses and actuation around one cycle; a separate study organizes synchronized responses in orbit–forcing space. Only their regular transporting components give local charts of the *Resonant Locomotion Manifold* (RLM). A fast loop locally regulates phase, while a slower amplitude correction regulates mean speed. The thesis thereby connects natural oscillation, conservative locomotion, resonant maintenance, and control within the stated local scope.

Keywords. mechanical locomotion; natural oscillations; nonlinear modes; geometric mechanics; relative-periodic motion; nonholonomic constraints; energy exchange; forced response; resonance; phase-locked loop.

Contents

Résumé	i
Abstract	ii
Notation and conceptual vocabulary	vii
I From natural oscillation to natural locomotion	1
1 Natural locomotion as a family-selection problem	2
1.1 Four insufficiencies determine the object	2
1.2 Conservative natural locomotion	3
1.2.1 From one cycle to a local family	4
1.3 Contributions and their logical status	5
1.4 Where the thesis enters established theory	7
1.5 Models, scope, and conceptual progression	8
2 From natural oscillation to natural locomotion	10
2.1 A natural oscillation is a family, not a sine wave	10
2.1.1 Closed energy loops	10
2.1.2 Action labels the loop; angle clocks the motion	11
2.1.3 The pendulum makes the amplitude dependence visible	12
2.1.4 What transfers to the locomotor problem	13
2.2 Why total energy no longer selects one motion	13
2.2.1 Linear modes are local seeds	14
2.2.2 Why the linear modal plane bends	15
2.2.3 Finite-amplitude modes change their choreography	15
2.2.4 Action and phase only after the modal family is chosen	16
2.3 The oscillator–locomotor analogies	17
2.4 Reduced periodicity leaves the pose open	17
2.4.1 Reduced dynamics and pose reconstruction	18
2.4.2 Exact reduced return and nontrivial group displacement	19
2.4.3 How planar pose increments are compared	19
2.4.4 A worked geometric loop: displacement without selection	20
2.4.5 Reconstruction does not certify dynamical compatibility	20
2.4.6 The common constrained-mechanics construction	21
2.5 From a returned cycle to a corrected family	22
2.5.1 An oriented section turns a cycle into a fixed point	22
2.5.2 Shooting corrects closure and a phase row fixes the clock	22
2.5.3 Pseudo-arclength follows the curve through a fold	23

2.5.4	When a computed family is locally a manifold	24
2.5.5	Candidate, corrected cycle, and mechanically qualified member	25
2.6	What losses change	26
2.6.1	From energy conservation to work balance	26
2.6.2	The forcing clock defines the return	26
2.6.3	Why quadrature is useful but not sufficient	27
2.7	The order of the construction	28
II	Natural locomotion under total-energy conservation	29
3	Natural-locomotion families of the two-segment locomotor	30
3.1	Mechanical model	30
3.1.1	Configuration, constraint, and pose reconstruction	30
3.1.2	Compact reduced dynamics	32
3.2	Straight motion and the oscillatory seed	33
3.3	Shooting and continuation	34
3.4	The continued family and its locomotor meaning	34
3.5	The scalar return coordinate	38
4	Effective oscillator, propulsive channel, and exact scalar return	40
4.1	The return coordinate to be replaced	40
4.2	Carrier elimination and the effective oscillator	41
4.2.1	Completed square and transverse storage	41
4.2.2	The closed transverse field	43
4.3	Closed-to-open lift and restoration of the full field	44
4.3.1	Restored finite-speed field	44
4.3.2	Finite-speed and orientation chart	45
4.4	Exact scalar exchange–return equivalence	45
4.5	A computable exchange residual	46
4.6	Numerical realization	47
4.6.1	Returned family	47
4.7	Dimensional boundary of the scalar result	47
5	Modal families of the three-segment locomotor	49
5.1	Why one storage value cannot identify the endpoint	49
5.2	Locomotor and reduced conservative field	50
5.2.1	Configuration, state, and pose reconstruction	50
5.2.2	Unforced equations and total energy	51
5.3	Closed supports, modal seeds, and the closed-to-open lift	52
5.3.1	Carrier-closed transverse storage	52
5.3.2	Local in-phase and anti-phase seeds	52
5.3.3	Closed-to-open lift and full-field correction	53
5.4	Opened cycles: vector return, modal classification, and exchange	53
5.4.1	Qualification by the full reduced field and vector return	53
5.4.2	Classification after return	54
5.4.3	Exchange after full return	54
5.4.4	Following one cycle from modal support to locomotion	54
5.5	Coexisting anti-phase- and in-phase-descended families	54
5.6	Finite-amplitude contrast between the sectors	56
5.7	Matched-straight pose decomposition	56

5.8	What transfers from the scalar construction	59
III	Dissipation, resonantly maintained locomotion, and local regulation	60
6	From passive families to resonantly maintained locomotion	61
6.1	From energy loss to synchronized return	61
6.2	A mechanics-preserving bridge for one conservative cycle	62
6.3	A separate parameter-distinct forced model	63
6.4	Synchronized-response sheets and their transported subset	65
6.4.1	Phase measured from deformation and torque	66
6.5	Speed and period across forcing-amplitude levels	66
6.6	Exact cycle work and a local quadrature marker	68
6.7	Pose reconstruction separates locomotion from oscillation	69
7	Local phase–speed regulation of resonant responses	71
7.1	Why phase is an ambiguous coordinate	72
7.2	Online phase and the locally tested frequency sign	72
7.3	Fast phase loop, slow speed loop, and guards	74
7.4	Regulating phase and speed across the rising interval	75
7.5	Sampled reduced-order evidence	76
7.6	Control result and scope	76
IV	Synthesis, limits, and outlook	77
8	Synthesis: accompanying intrinsic locomotor families	78
8.1	The scientific result	78
8.2	One constrained-mechanics method, two locomotors	78
8.3	Families, transport, and local manifold geometry	79
8.4	Energy exchange: exact scalar closure and its vector boundary	80
8.5	From passive persistence to resonant maintenance	80
8.6	What local regulation adds	81
8.7	Limits that define the next questions	82
8.8	Conclusion	82
A	A common reduction for continuously constrained locomotors	84
A.1	Admissible velocities and virtual displacements	84
A.2	Projected Lagrange–d’Alembert equation	84
A.3	Why ideal constraints redistribute but do not energize	85
A.4	The two model instances	85
A.5	Scope of the shared method	85
B	Two-segment model and conventions	86
B.1	Geometry, velocities, and ideal constraint	86
B.2	Dimensional reduced equations	86
B.3	Dimensionless dictionary and energy check	87
B.4	Population-level numerical checks	87
B.5	Conventions carried into the numerical chapters	88
C	Scalar field, proof details, and numerical protocol	89

C.1	Coefficient functions and opened field	89
C.2	Exchange identity in the implemented convention	90
C.3	Event integration and nonlinear solve	90
C.4	Regular domain and numerical validation	91
D	Three-segment model and returned-cycle population	93
D.1	Exact mechanical model	93
D.2	One numerical model-parameter vector	94
D.3	Modal coordinates and numerical construction	95
D.4	Cycle qualification and modal identity	95
D.5	Population verification	96
D.6	What the construction tests separate	96
D.7	Observable convention and scope	97
E	Forced-response equations and numerical verification	98
E.1	Numerical construction of the mechanics-preserving bridge	98
E.2	Equations of the parameter-distinct forced model	98
E.2.1	Population-wide dynamic-return audit	99
E.3	Phase extracted from the continued cycles	100
E.4	Peak-speed and quadrature extraction	100
E.5	Population-wide pose reconstruction	101
E.6	Nonlinear shift from the small-amplitude period	101
F	Phase-speed controller and sampled finite-harmonic evidence	102
F.1	Rolling phase and speed estimates	102
F.2	Sampled phase and amplitude laws	102
F.3	Sign calibration and representative amplitude levels	103
F.4	Stage-transition logic and terminal metrics	103
F.5	Terminal-cycle, time-step, and rejection audits	104
F.6	Closed finite-harmonic controller matrices	105
F.7	Requirements for a finite-arrival theorem	105
G	Local geometry of continued locomotion families	106
G.1	Removing the arbitrary time shift	106
G.2	From a phase-fixed branch to an invariant family	106
G.3	A manifold and its observable image are not the same object	107
G.4	Natural and resonant locomotion objects	107
G.5	What a numerical regularity check would establish	107

Notation and conceptual vocabulary

The same symbols recur across several mechanical models. These two tables are a consultation aid, not a list of prerequisites: every concept is introduced progressively in the text before it is used in a derivation. Model-specific states, coefficients, and sign conventions are defined where they first appear.

Symbol or term	Meaning
$g = (x, y, \theta) \in \text{SE}(2)$	planar body pose; it may accumulate a nonzero net translation or rotation while the reduced state returns
e_G	identity element of the pose group
$r, \delta, \text{ or } \boldsymbol{\delta}$	internal shape coordinate, one joint angle, or a vector of joint angles
z	pose-reduced mechanical state, including the rates needed to advance the dynamics
v, ω, σ	body-forward speed, body yaw rate, and joint rate
$\bar{v}$	cycle-mean body-forward speed
Δg	pose increment obtained by pose reconstruction after one reduced return
Δg_{inc}	matched-straight residual transform; it is not a causal energy partition
Σ	oriented Poincaré section used to select one event on a cycle
$E_{\text{prop}}, E_{\perp}$	propulsive-channel and effective-transverse-oscillator storage
$P_{\text{ex}}, I_{\text{ex}}$	exchange power and its time integral, with the sign declared in Chapter 4
$A_{\text{exc}}, T_{\text{exc}}$	forcing amplitude and forcing period
$\phi_{\delta\tau}$	fundamental phase of joint deformation relative to applied joint torque

Concept	Meaning in this thesis
Two-segment locomotor (2SEG)	the articulated locomotor with one effective transverse oscillatory degree of freedom
Three-segment locomotor (3SEG)	the articulated locomotor with two effective transverse oscillatory degrees of freedom
In-phase (IP) / anti-phase (AP)	linear modal seeds; finite-amplitude members are described as continued from these seeds unless exact phase locking is established
Cycle	one periodic or relative-periodic solution, considered up to time shift
Phase-fixed branch	a locally one-dimensional zero set of independent return equations after the arbitrary cycle phase has been removed
Local manifold chart	a parameterized neighbourhood of a regular solution object: one-dimensional after phase fixing and locally two-dimensional after cycle phase is restored, under the stated rank and embedding hypotheses
Effective transverse oscillator	the oscillatory subsystem exposed by carrier elimination and described by transverse storage $E_{\perp}$
Propulsive carrier channel	the complementary admissible carrier motion; shortened to <i>propulsive channel</i> in physical explanations
Carrier elimination / reopening	construction of an auxiliary carrier-closed field / restoration of the full opened mechanical field
Pose reconstruction	integration of the group equation after the reduced motion is known; it is distinct from carrier reopening
Natural Locomotion Manifold (NLM)	the regular local solution object associated with exact unforced reduced return and open pose; computed components are reported as families or branches
Synchronized-response family ($\mathcal{S}_R$)	the solution set of the one-to-one forced return equations before transport and local regularity are inferred
Resonant Locomotion Manifold (RLM)	regular local components of the transported subset $\{s \in \mathcal{S}_R : \Delta g(s) \neq e_G\}$; plotted phase, speed, waveform, and pose quantities are observable projections
Phase-locked loop (PLL)	the fast phase-feedback loop used within local phase-speed regulation over the tested observable interval

Throughout the thesis, *returned* refers to the complete reduced mechanical state specified by the model, not merely to shape. Locomotion also requires $\Delta g \neq e_G$. The word *manifold* is local unless a stronger statement is made explicitly; it implies neither attraction, stability, optimality, uniqueness, nor global smoothness.

Part I

From natural oscillation to natural locomotion

Chapter 1

Natural locomotion as a family-selection problem

A locomotor converts internal deformation into translation or rotation of its body through mechanical interaction with the environment. It can be commanded by prescribing every joint trajectory. That strategy is appropriate when the desired choreography is already known, but it does not reveal which motions the mechanism can organize through its passive dynamics. Elasticity, inertia, constraints, and morphology act before a controller is chosen. The question of this thesis is therefore prior to trajectory tracking: *which repeatable motions are supported by the locomotor's intrinsic dynamics—its declared passive mechanical field—and how can they later be maintained without overwriting that organization?*

The answer cannot be one privileged sinusoid. Even a pendulum changes its period and waveform with amplitude. An articulated locomotor adds two further features: several internal coordinates can reorganize their relative timing, and the body must advance while the internal mechanical event repeats. The object is consequently a family of timed motions, not a geometric path or a fixed joint template.

Figure 1.1 gives the whole thesis in one view. Prescribed choreography starts from a time history and asks the mechanics to follow it. Natural-family construction starts from the unforced equations and searches for exact reduced return with open pose. Control is introduced only after losses, where it regulates phase and mean-speed observables instead of prescribing every harmonic.

1.1 Four insufficiencies determine the object

Four successive insufficiencies lead from a familiar natural oscillation to the locomotor object studied here. Each step retains what the preceding theory provides and makes explicit what remains missing.

A natural oscillation is not one fixed sine wave. Near an equilibrium, an eigenvector gives a modal direction and its eigenvalue gives a small-amplitude natural frequency. At finite amplitude, nonlinear forces can bend the orbit, change its period, create harmonics, and alter the relative phase and amplitude of its coordinates. Simply scaling the eigenvector therefore imposes a choreography that need not solve the nonlinear equations. Nonlinear modal families supply the required finite-amplitude organization.

A full-state oscillation cannot produce net transport. Periodicity always means closure in a specified state space. An ordinary periodic orbit returns every component of configuration and velocity. If body position and orientation belong to that state, they return too, so the net rigid-body displacement is the identity. A locomotor requires a different closure: its reduced mechanical state

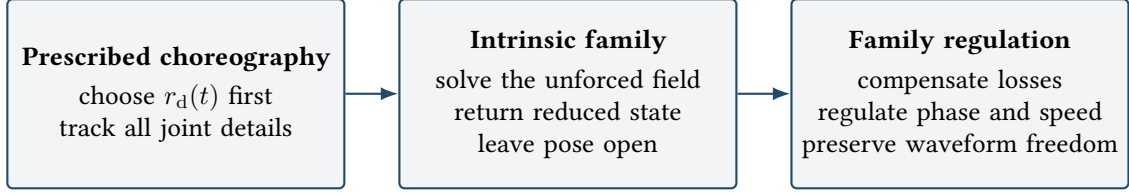


Figure 1.1: Three distinct design viewpoints. The thesis first replaces an imposed joint history by a family selected through the conservative mechanics. It then studies how related dissipative responses can be maintained and locally regulated through phase and mean speed.

must form a closed loop while its pose is allowed to accumulate a net translation or rotation. The relevant motion is *relative-periodic*, periodic only after quotienting pose symmetry.

A transporting internal loop need not be natural. Locomotion geometry, formalized by geometric mechanics, can reconstruct a body displacement from a given loop of internal shapes. This answers how a loop moves the body, but not why that loop should occur. Infinitely many loops can be drawn, optimized, or imposed by actuators, and their timing matters in an inertial mechanism. A natural candidate must be a time-parameterized solution of the passive, unforced mechanical field, not merely a closed curve with a favourable displacement.

Conservative return does not survive losses without input. An exactly initialized member of an invariant conservative family persists in the ideal model because no energy is removed. This is invariance, not attraction or robustness. Once damping is introduced, a free oscillation decays; a periodic response can persist only if input work compensates the dissipated work. The forced problem therefore has its own return equation and its own solution family. The controller studied later regulates phase and mean speed over a tested observable interval rather than enforcing a precomputed joint waveform.

The resulting chain is concise:

$$\begin{aligned}
 &\text{fixed choreography} \longrightarrow \text{finite-amplitude intrinsic family,} \\
 &\text{full-state periodicity} \longrightarrow \text{reduced return with open pose,} \\
 &\text{transporting loop} \longrightarrow \text{dynamically compatible loop,} \\
 &\text{conservative family} \longrightarrow \text{loss-compensated response and local regulation.}
 \end{aligned} \tag{1.1}$$

1.2 Conservative natural locomotion

Let the configuration split into planar pose and internal shape,

$$q = (g, r), \quad g \in \text{SE}(2), \tag{1.2}$$

where $\text{SE}(2)$ is the group of rigid translations and rotations in the plane, and let z denote the pose-reduced mechanical state. Besides shape, z contains its rates and admissible *carrier* velocities: the forward and yaw velocities of the main body segment used below. Pose symmetry yields the reduced and reconstruction equations

$$\dot{z} = f(z), \quad \dot{g} = g \widehat{\xi(z)}. \tag{1.3}$$

The first equation decides which internal motions are mechanically possible. The second integrates the body pose after such a motion has been found.

A conservative natural-locomotion cycle is a nontrivial solution satisfying

$$\dot{z} = f_{\text{cons}}(z), \quad z(T) = z(0), \quad \Delta g := g(0)^{-1}g(T) \neq e_G. \tag{1.4}$$




Figure 1.2: The repeatable locomotor event. The reduced state contains the internal coordinates, their rates, and the required carrier velocities; it returns after one cycle. Pose reconstruction remains open and accumulates a nonidentity group displacement.

The first condition excludes an arbitrary choreography; the second makes the same reduced mechanical event repeat; the third distinguishes locomotion from oscillation in place. A straight relative equilibrium, for which z is constant while pose advances at constant velocity, is a seed or boundary of the computed families. It is not converted into a nontrivial periodic orbit by assigning it an arbitrary period. Throughout this thesis, *natural* has the narrow operational meaning of compatibility with the declared unforced mechanics. Preference, stability, and optimality are separate properties and are stated only when they are established.

Figure 1.2 shows why pose is deliberately omitted from the return equation. The reduced loop closes, but reconstruction composes a nonidentity rigid-body increment. Repeating the same reduced cycle repeats this body-frame increment.

The nonidentity increment is a locomotor output, but it is not yet a causal partition of propulsion. A cycle can modulate a large net translation already present at the straight relative equilibrium. Whenever an oscillation-induced increment is discussed below, the comparison is made in $SE(2)$ against a matched straight motion; the resulting residual transform is reference dependent.

The concrete anchor used throughout the thesis is the *two-segment locomotor* (2SEG): a carrier bearing the continuously active no-side-slip constraint and a second segment connected through an elastic joint. Its reduced state is $z = (\delta, v, \omega, \sigma)$, comprising joint angle, forward carrier speed, carrier yaw rate, and joint rate. Figure 1.3 shows one exported continuation cycle before the detailed model is derived. The left panel shows the coincident stored endpoints of the reduced state; the right panel reconstructs the carrier pose and makes the open endpoint visible. The example therefore visualizes the two properties imposed in Equation (1.4): internal repetition and external transport. Independent exact-field re-integration and local-rank evidence are reported later for 29 matched members, rather than inferred from this plot alone.

1.2.1 From one cycle to a local family

One isolated cycle is not yet a manifold. The intended family contains nearby cycles obtained by varying one mechanical label while correcting the complete return conditions. Each cycle also carries an arbitrary starting time, so the same physical orbit would otherwise be counted infinitely many times. A section or phase condition removes that redundancy, and the independent return equations may then be written schematically as

$$F(u, \lambda) = 0, \quad F : \mathbb{R}^{N+1} \rightarrow \mathbb{R}^N, \quad (1.5)$$

where u represents the phase-fixed orbit and λ is a continuation coordinate. If F is smooth and its Jacobian has rank N , the implicit function theorem makes the nearby zero set a smooth one-




Figure 1.3: One returned cycle of the two-segment locomotor. (a) The joint angle–rate projection closes; the stored full reduced state $(\delta, v, \omega, \sigma)$ returns at the same endpoint. (b) Integrating the carrier velocity and yaw rate reconstructs distinct initial and final poses.

dimensional branch. A fold in a plotted parameter does not destroy this local curve; it only makes that parameter a poor local coordinate. Restoring phase along every nontrivial cycle sweeps a locally two-dimensional family in reduced state space, provided the cycles remain embedded and belong to one fixed vector field.

For regular embedded branches whose reconstructed pose is nonidentity, this is the precise local sense in which the thesis retains the term *Natural Locomotion Manifold* (NLM). Computed components are called families or branches because global smoothness is not established across modal sectors, bifurcations, different mechanisms, or different state dimensions. The word *manifold* adds no claim of stability, attraction, uniqueness, or optimality. Chapter G gives the phase-fixing, regular-level-set, and observable-map statements behind this terminology.

1.3 Contributions and their logical status

The contributions progress from conservative construction to forced regulation. Their analytical centre is the scalar exchange–return theorem; numerical family constructions and a local control demonstration complete the progression. Stating them here separates these results from the classical tools used to obtain them, as summarized in Table 1.1.

Table 1.1: Contribution sequence and evidence boundary. Standard modal, geometric, continuation, forced-response, and feedback tools are identified separately in [Section 1.4](#); the entries here are their thesis-specific mechanical use and result.

Contribution	Result	Status and boundary
Two-segment natural family	A transverse spectral seed is corrected and continued into finite-amplitude moving relative-periodic cycles whose waveform, coordination, and pose are reconstructed.	Model-specific numerical construction. The 379 exported records support population observables; independent return and rank checks cover 29 matched roots. A matched SE(2) reference shows that straight motion generally accounts for most of the total translation.
Effective oscillator, propulsive channel, and scalar theorem	Carrier elimination exposes transverse storage. After reopening the propulsive channel, integrated energy exchange is equivalent to the last unresolved scalar return coordinate under the stated section hypotheses.	Proved necessity-and-sufficiency theorem, conditional on finite speed, transversality, orientation, and storage injectivity. It is not a general existence, regularity, or stability theorem.
Three-segment modal families	Full vector return and independent modal metrics construct and distinguish moving families continued from in-phase and anti-phase seeds at one fixed parameter vector.	Numerical construction and classification. Scalar storage is not sufficient for return; global modal genealogy and stability are not established.
Resonantly maintained locomotion	A mechanics-preserving local continuation connects one conservative cycle to damped, forced responses. A separate parameter-distinct study organizes synchronized locomotor responses in orbit-forcing space.	One local bridge and one independent response family. The work-loss balance is exact on periodic responses; quadrature is a local observable, not a universal selector.
Local phase-speed regulation	Fast phase feedback and slower mean-speed feedback regulate the two observables across a tested rising interval.	Nominal local numerical demonstration with sampled reduced linearizations; no online membership test for a stored response, global acquisition, exact nonlinear stability, robustness, or hardware claim.

The central analytical step deserves one additional sentence. The admissible forward and yaw motion of the carrier forms the propulsive channel through which the body translates and turns. Ideal constraint reactions perform no total work, but the admissible velocity distribution and kinetic coupling organize how total energy is shared between an effective transverse oscillator and this channel. In the two-segment locomotor, closure of that exchange identifies exactly one missing scalar return. In the three-segment locomotor, the corresponding transverse endpoint is vector valued, so full return must be solved before exchange or modal quantities are interpreted. This dimensional boundary is as important as the scalar theorem itself.

The forced solution object is defined with the same care. The one-to-one forced-return equations define a synchronized-response set in orbit-forcing space. Pose reconstruction selects its transporting members; regular local components of that transported subset are referred to as the Resonant Locomotion Manifold (RLM). Frequency, phase, waveform, mean speed, and reconstructed pose are observable coordinates or projections, not the manifold by definition. The mechanics-preserving local bridge and the parameter-distinct response family remain separate constructions.

1.4 Where the thesis enters established theory

The literature is most useful when organized by the question it answers rather than by a list of communities.

Why use passive mechanics as a motion resource? Passive-dynamic walking established that a mechanism can organize a repeatable gait with little or no direct trajectory enforcement, and physical robots later showed how passive structure can be combined with modest actuation (McGeer, 1990; Collins et al., 2005). Reviews of energy-efficient locomotion and resonance-exploiting mobile robots place this idea in a broader design and control landscape (Kashiri et al., 2018; Schönebaum et al., 2021). Biological mechanics offers both motivation and a caution: body resonance can reduce cyclic effort, but fluid loading, several resonant peaks, and aerodynamic effects prevent resonance from being a universal explanation of locomotor performance (Ahlborn et al., 2006; Tytell et al., 2014; Ramanananarivo et al., 2011). The thesis therefore uses *natural* as a model-relative dynamical criterion, not as a synonym for efficient or biologically optimal.

Which free motions form coherent finite-amplitude families? Linear modes provide tangent seeds, while nonlinear normal modes organize periodic families whose frequency and waveform vary with amplitude (Rosenberg, 1966; Shaw and Pierre, 1993; Kerschen et al., 2009; Peeters et al., 2009; Albu-Schäffer and Della Santina, 2020). In robotics, nonlinear modal surfaces have also been used as feedback targets, first in control constructions and then in closed-loop experiments on an elastic industrial robot (Della Santina and Albu-Schäffer, 2021; Bjelonic et al., 2022). Conservative and passive locomotion studies likewise demonstrate continua of gaits and transitions organized by energy, design, and bifurcation structure (Remy, 2011; Gan et al., 2018; Raff et al., 2022). Natural-motion-manifold control shows how passive families can serve as control targets (Calzolari et al., 2025). These works establish that families exist and can matter for control. The present thesis adds an exact reduced-return construction and its dimensional interpretation for the declared articulated locomotors.

How can internal repetition coexist with body displacement? Geometric and nonholonomic mechanics separate shape from pose and reconstruct rigid-body motion from reduced variables (Kelly and Murray, 1995; Ostrowski and Burdick, 1998; Hatton and Choset, 2011; Bloch, 2015). Reduced models with passive internal degrees of freedom and internal-rotor systems further show how unactuated shape dynamics, nonholonomic constraints, and locomotion interact (Boyer and Belkhir, 2014; Tallapragada, 2015; Fedonyuk and Tallapragada, 2019). Relative-periodic orbits express return modulo a symmetry action (Fasso et al., 2020; Eldering and Jacobs, 2016). These tools supply the open-pose language used here. They do not select which time-parameterized internal loop solves the passive mechanics; that selection is supplied by the reduced field and return residual.

How are cycles and branches computed? Shooting and pseudo-arclength continuation correct periodic solutions and follow residual-defined branches through folds (Peeters et al., 2009; Kuznetsov, 2004). The numerical method is deliberately agnostic: it follows whichever equations it is given. The contribution is therefore not continuation itself, but the mechanical construction of the independent residuals, including the scalar exchange–return equivalence and the full vector return retained for the three-segment case.

How are related motions maintained when energy is lost? Backbone extraction, cycle work balance, forced response, phase resonance, and nonlinear testing provide the appropriate dissipative language (Cenedese and Haller, 2020a,b; Peeters et al., 2011; Kuether et al., 2015). Phase-locked and control-based continuation methods can track nonlinear responses in numerical and experimental

settings (Peter et al., 2016; Peter and Leine, 2017; Denis et al., 2018; Renson et al., 2016; Barton, 2017; Abeloos et al., 2022). The forced contribution here combines a locomotor response family with reconstructed pose and mean speed, then regulates phase and speed locally. It does not claim that phase-locked-loop continuation is itself new.

Related results have appeared in co-authored work on the two-segment family, forced response, and the closed–open construction (Rajaomarasata et al., 2025a,b; Mortel et al., 2026). The dissertation places them in one mechanical notation and makes explicit the progression from scalar exchange closure to vector return and then to forced local regulation.

1.5 Models, scope, and conceptual progression

The two mechanisms studied in detail are normalized planar articulated locomotors with smooth, bilateral, ideal, continuously active no-side-slip constraints of constant rank. The conservative models contain elastic storage but no damping or maintained actuator input. The local forced bridge retains one conservative model’s configuration, constraint distribution, inertia, and potential while adding a declared loss law and periodic joint torque. The larger forced-response study retains the two-segment kinematic architecture but uses a different parameter vector and elastic law; it is not a continuation of the same conservative family.

The common constrained-mechanics method is broader than these two examples. The velocities allowed by the constraints form the kernel of the constraint matrix. A basis of this admissible space parameterizes velocity, the Lagrange–d’Alembert equations are projected onto it, the reduced field is integrated, and pose is reconstructed. Its stated domain does not include impacts, unilateral contact, changing contact modes, stick–slip, or fluid memory. Hybrid walking and fluid–structure locomotion are therefore research extensions, not cases already covered by the results.

Availability and reproducibility boundary. The accompanying repository contains the manuscript source, compact executable model snapshots, the promoted two-segment cycle surface, audit scripts, and machine-readable audit outputs. It does not bundle the full historical two-segment continuation logs or the complete raw three-segment and forced-response cycle archives, and no public DOI or durable download location is claimed for those heavy inputs. Recorded provenance and input hashes distinguish re-computable post-processing from audits whose source archives remain external. A clean checkout can therefore build the manuscript and rerun the calculations whose inputs are included, but it cannot regenerate every continuation population from zero.

The frozen scientific release is identified by the Git tag `thesis-v4-scientific-release` in the private GitHub repository at github.com/rajaommi/thesis_nlm_workspace; access requires repository authorization. The tagged checkout currently declares no general reuse license, so access is not permission to redistribute text, figures, code, or data. The source archive for the conservative-to-forced anchor and the external 3SEG and forced archives presently have no public DOI; their durable institutional deposit and reuse terms remain publication tasks. Table 1.2 separates what can be rerun from the tagged checkout from what additionally requires an external source archive. Exact copy-paste commands and full file inventories are retained in the repository evidence ledger `rewrite/evidence/REPRODUCIBILITY.md`.

The argument can be followed through three conceptual stages, summarized in Table 1.3. The appendices carry coefficient algebra and reproducibility detail, but the concepts and inferences required to understand the contributions remain in the main text.

Chapter 2 supplies the minimum cross-community foundations needed for these milestones. It starts from a pendulum because one degree of freedom makes family, action, and phase visible. It then uses the double pendulum to explain why total energy is no longer selective, before introducing open pose, exact reduced return, continuation, local family geometry, and the common constrained modelling pattern. Each idea is carried only to the depth needed by the next result.

Table 1.2: Reproduction routes and access boundary. Script names are abbreviated here; exact root-relative invocations and inventories are given in the evidence ledger. Audit outputs record SHA-256 identifiers, settings, and environment. The verified run used Python 3.12.10, NumPy 2.3.4, and SciPy 1.16.2.

Result	Included in the tagged checkout	Additional input and access	Portable command / output
2SEG return and rank	exact model, 29 section states, comparison rows	none	audit_2seg_return_rank.py → 2seg_return_rank_audit.json
Same-model forced bridge	model, continuation script, derived branch and audits	9,354,125-byte 2SEG anchor archive, external; pass -two-seg PATH	analyze_forced_bridge_2seg.py → gate_b_local_bridge_summary.json
3SEG dynamic return	exact model and verifier	11-MB summary and referenced cycles, external; set NLM_3SEG_SUMMARY	audit_3seg_dynamic_return.py -summary \$NLM_3SEG_SUMMARY → 3seg_dynamic_return_audit.json
Forced response	exact model, from-zero producer, compact 25-file hash manifest	166,574,295-byte raw archive, external; set NLM_OCEANS_RAW_DIR	audit_oceans_dynamic_return.py -raw-dir \$NLM_OCEANS_RAW_DIR -verify-symbolic-model → oceans_dynamic_return_audit.json
Phase-speed control	model, controller, 25 reference cycles, stage table	none for the reported follow-up checks	generate_guarded_follower.py -audit-control -skip-figure → v4_control_followup_audit.json

Table 1.3: Conceptual progression and comprehension milestones.

Conceptual stage	Chapters	Milestone
Object and construction	Chapters 1–3	Distinguish nonlinear family from fixed sinusoid, reduced return from full-state periodicity, and pose reconstruction from dynamical compatibility; follow the first two-segment family from seed to pose.
Conservative qualification	Chapters 4–5	Explain the effective oscillator and propulsive channel, the scalar exchange–return theorem, its section hypotheses, and why the three-segment locomotor still requires vector return and modal identity.
Losses and regulation	Chapters 6–8	Explain input–loss balance, the local conservative-to-forced bridge, the distinction between a synchronized-response sheet, its transported subset, the regular local RLM components, and their observable projection, together with the local phase–speed control result.

Chapter 2

From natural oscillation to natural locomotion

The central object of the thesis combines ideas that are elementary in their home disciplines but easy to conflate when placed in one locomotor model. The chapter therefore rebuilds them from the smallest physical examples and introduces each mathematical term only when it answers a concrete question. A frequency can label a nonlinear family without being constant across that family. A periodic internal motion can transport a body only if pose is left out of the return equation. A reconstructed displacement says what a loop does, not whether the passive dynamics generates it. A smooth continuation plot can be a projection of a regular solution branch without proving one global manifold. Finally, a periodic response can survive damping only through input work.

This chapter develops these statements in the order in which they become necessary. Each foundation answers a concrete problem, is illustrated on the smallest useful example, and ends at the point where the thesis-specific construction takes over.

Here the *carrier* is the main body segment. Eliminating its forward and yaw velocities temporarily exposes the effective transverse oscillator; reopening the propulsive channel restores them, whereas pose reconstruction later integrates position and orientation.

2.1 A natural oscillation is a family, not a sine wave

A small-amplitude mass–spring oscillator suggests a familiar picture of natural motion: one sinusoid with one natural frequency. The problem with this picture is not that it is false, but that it is local. A pendulum released through a large angle is still moving freely under its own conservative mechanics, yet it slows near its turning points, moves fastest through the bottom of the well, and takes longer to complete a cycle as the release angle increases. Periodicity survives; a fixed sinusoidal waveform and a frequency independent of amplitude do not. The first foundation must therefore identify an entire family and provide separate coordinates for *which member* is moving and *where on that member* the state lies.

2.1.1 Closed energy loops

Consider one local oscillator coordinate x , with momentum $p = M(x)\dot{x}$, positive inertia $M(x)$, and potential $U(x)$. Its Hamiltonian is

$$H(x, p) = \frac{p^2}{2M(x)} + U(x) = E. \quad (2.1)$$

Here H is simply the total mechanical energy written in position–momentum coordinates: kinetic energy plus potential energy. The Hamiltonian vocabulary will be useful below, but no prior




Figure 2.1: Energy organizes a family of bounded pendulum oscillations. Left: each energy below the separatrix intersects the potential at two turning points. Right: the corresponding momentum branches form a closed phase-space loop. Increasing the energy changes both the loop and its period. The q in this pedagogical figure is a local pendulum coordinate, not the full locomotor configuration used later.

Hamiltonian formalism is required for this one-degree-of-freedom picture. Energy conservation confines a free trajectory to a level set of H . In a regular potential well, an energy below the escape level intersects the potential at two turning points. The positive- and negative-momentum branches join there to form a closed curve in the phase plane. One circuit of that curve is one natural oscillation.

Figure 2.1 makes this construction visible. The potential selects the turning points, while the two possible signs of momentum complete the state-space loop; a displacement interval alone would omit half of the mechanical state.

The family, rather than one waveform, is the useful object. Each regular energy value selects another loop, and the time needed to traverse the loop is

$$T(E) = 2 \int_{x_-(E)}^{x_+(E)} \sqrt{\frac{M(x)}{2\{E - U(x)\}}} dx, \quad \Omega(E) = \frac{2\pi}{T(E)}. \quad (2.2)$$

The integrand becomes large near an ordinary turning point because the velocity vanishes there. This unequal time of flight is precisely why the physical coordinate is generally not a cosine. With constant inertia, a quadratic potential is the special case for which the period is independent of E and the waveform is sinusoidal.

2.1.2 Action labels the loop; angle clocks the motion

Energy is a physical family label. Action provides a canonical geometric label for the same loop:

$$J(E) = \frac{1}{2\pi} \oint_{H=E} p dx. \quad (2.3)$$

The contour integral is the oriented area enclosed in the phase plane divided by 2π . It measures the size of the whole state-space cycle, rather than only the largest displacement. On a regular annulus of such loops, E can locally be written as $H(J)$. An angle φ can then be chosen as a uniform clock along each member:

$$\dot{J} = 0, \quad \dot{\varphi} = \Omega(J) = \frac{dH}{dJ} = \frac{2\pi}{T(J)}. \quad (2.4)$$

Thus J says which oscillation is taking place and φ says where the state lies on it. These are the Hamiltonian analogues of radius and polar angle, but the analogy must not be pushed too far: the radius-like quantity is a phase-space area, and the physical loop need not be circular (Arnold, 1989).




Figure 2.2: The same finite-amplitude pendulum libration viewed in physical configuration and in action–angle phase. The pendulum spends more time near its turning points, whereas the angle variable advances uniformly around the abstract circle. Uniform phase therefore coexists with a non-sinusoidal physical waveform.

Calling the pair canonical means that it preserves the mechanical phase-space area,

$$dx \wedge dp = d\varphi \wedge dJ. \quad (2.5)$$

This is what distinguishes action–angle coordinates from an arbitrary redrawing of a closed curve as a circle.

Uniform phase also does not make the measured coordinate sinusoidal. The inverse map has the form

$$x(t) = X(J_0, \varphi_0 + \Omega(J_0)t), \quad (2.6)$$

where $X(J, \varphi)$ is periodic in φ but may contain several harmonics. The nonlinearity has been moved into this map; it has not been removed.

Figure 2.2 contrasts the physical and intrinsic clocks. Equal time intervals are unevenly spaced along the pendulum arc but evenly spaced in the action–angle phase. The angle is therefore a dynamical clock, not the geometric polar angle of a distorted phase-plane loop.

2.1.3 The pendulum makes the amplitude dependence visible

Let ϑ be the pendulum angle, I its inertia, and $\omega_0 = \sqrt{g/l}$ its small-amplitude frequency. The exact Hamiltonian is

$$H(\vartheta, p_\vartheta) = \frac{p_\vartheta^2}{2I} + I\omega_0^2(1 - \cos \vartheta). \quad (2.7)$$

For a libration of maximum angle $a < \pi$, set $k = \sin(a/2)$. Its period and frequency are

$$T(a) = \frac{4}{\omega_0} K(k), \quad \Omega(a) = \frac{\pi\omega_0}{2K(k)}, \quad (2.8)$$

where

$$K(k) = \int_0^{\pi/2} \frac{d\psi}{\sqrt{1 - k^2 \sin^2 \psi}} \quad (2.9)$$

is the complete elliptic integral of the first kind. Near the stable equilibrium, Ω tends to ω_0 . As a approaches the upright separatrix, the period diverges and the frequency tends to zero.

The same mechanism therefore carries a continuum of natural frequencies, one for each member of the libration family.




Figure 2.3: Finite-amplitude organization of pendulum librations. The period grows and the frequency decreases with maximum angle, while the action grows with the area enclosed by the phase-space loop. The dashed boundary is the separatrix limit; it is not part of the regular action–angle annulus.

Figure 2.3 places the two family coordinates side by side: action grows with the area of the loop, while frequency decreases as the separatrix is approached. A single number ω_0 describes only the common small-amplitude limit.

This construction is deliberately local. It applies on an annulus filled by regular compact energy loops. The action–angle chart is singular at the equilibrium and at the separatrix and does not provide one coordinate system across both librations and rotations. Nor does the one-degree argument imply global action–angle coordinates for a general multi-degree system. Its useful message is narrower: a nonlinear natural oscillation has a family label, an orbit-dependent frequency, and a uniform phase even when its physical choreography is not sinusoidal.

2.1.4 What transfers to the locomotor problem

The downstream role of this example is structural. Carrier elimination in the two-segment locomotor will expose a transverse storage of the form

$$E_{\perp}(\delta, \sigma) = \frac{1}{2} M_{\perp}(\delta) \sigma^2 + U(\delta), \quad (2.10)$$

where δ is the internal angle and $\sigma = \dot{\delta}$. On a regular bounded component, the storage level plays the same organizing role as the pendulum energy loop: it identifies a one-degree transverse oscillation, its period, and an action-like family coordinate. The analogy explains why the two-segment case admits a scalar return argument.

It does not yet produce locomotion. The pendulum has no propulsive channel and no pose to leave open. In the locomotor, the carrier-closed oscillation is only an auxiliary support; the full propulsive channel must be restored, the opened reduced state must return, and pose must then be reconstructed. This precise limit keeps the pendulum analogy useful without equating the two mechanisms.

With several oscillatory degrees of freedom, energy no longer identifies one closed loop. A further organizing object is then required before this one-family-label/one-phase picture can be recovered locally.

2.2 Why total energy no longer selects one motion

The scalar construction works because a regular energy level in a two-dimensional phase plane can itself be a closed orbit. For a conservative system with n degrees of freedom, the phase space has

dimension $2n$, and a regular energy level has dimension

$$\dim H^{-1}(E) = 2n - 1. \quad (2.11)$$

With two degrees of freedom, one energy value therefore leaves a three-dimensional surface. That surface may contain distinct periodic motions, quasi-periodic trajectories, exchanges of energy between coordinates, and resonant or irregular dynamics. Conservation of energy remains necessary, but it does not select one coherent oscillation. In particular, not every free trajectory on an energy surface is a mode.

This is the exact obstruction that the double pendulum is meant to expose. In the simple pendulum, the intersection $H = E$ is already the one-dimensional orbit. In the double pendulum, $H = E$ is three dimensional, whereas one periodic orbit is still one dimensional. Two additional pieces of selection are missing. A modal family supplies them by choosing a coherent two-dimensional modal surface whose intersection with a regular energy level is a periodic member.

2.2.1 Linear modes are local seeds

Near a stable equilibrium, let $\mathbf{u}$ denote a small displacement. The linearized equations are

$$M_0 \ddot{\mathbf{u}} + K_0 \mathbf{u} = 0, \quad K_0 \phi_j = \omega_j^2 M_0 \phi_j. \quad (2.12)$$

Each eigenvector ϕ_j gives a configuration-space direction, and the corresponding displacement-velocity pair spans a two-dimensional modal plane in the linearized state space. Motion initialized in that plane is sinusoidal and remains there under the linear equations.

The symmetric double pendulum supplies a concrete picture. Its small-angle model has a lower-frequency in-phase seed, for which the link angles have the same sign, and a higher-frequency anti-phase seed, for which they have opposite signs. In the nondimensional example shown in Figure 2.4, the modal directions are proportional to $(1, \sqrt{2})^\top$ and $(1, -\sqrt{2})^\top$, respectively.




Figure 2.4: The two small-angle modal seeds of the symmetric double pendulum. The in-phase direction moves both links to the same side; the anti-phase direction moves them to opposite sides. These are tangent directions at the downward equilibrium, not finite-amplitude trajectories. The angle symbols in the source figure are local to this example.

The word *local* is essential. A linear eigenvector describes the first variation of the dynamics at one equilibrium. Increasing its amplitude while keeping the same proportional sine waves does not generally solve the nonlinear equations.

2.2.2 Why the linear modal plane bends

Write the nonlinear equations in linear modal coordinates η_j :

$$\ddot{\eta}_j + \omega_j^2 \eta_j = N_j(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}}), \quad (2.13)$$

where N_j collects the nonlinear inertia, stiffness, and coupling terms. In the linear model the right-hand side vanishes. In the nonlinear model, setting all modes except mode k to zero need not make $N_j = 0$ for $j \neq k$. The vector field then points out of the flat linear modal plane. The plane is a correct tangent approximation, but it is not an exact finite-amplitude invariant support.

A nonlinear modal family corrects this failure. Starting from a small orbit near one linear seed, the full nonlinear equations determine the changes in period, relative amplitudes, phase relations, and harmonics required for exact closure. Continuing the corrected orbit produces a family rather than one amplified eigenvector. In practical conservative modal analysis, this family of periodic motions is a standard representation of a nonlinear normal mode (Rosenberg, 1966; Shaw and Pierre, 1993; Kerschen et al., 2009; Peeters et al., 2009; Albu-Schäffer and Della Santina, 2020).

Under suitable local regularity and non-resonance assumptions, the same organization can be described geometrically by a two-dimensional invariant surface tangent to the selected linear modal plane and filled locally by periodic orbits. Existence, uniqueness, and smoothness of such invariant objects require hypotheses; they do not follow merely from plotting a continued branch (Haller and Ponsioen, 2016). The term *family* is used below whenever only the periodic-orbit organization is needed.

Locally, one may picture such a surface through an embedding

$$W_j : (a, \varphi) \mapsto (\mathbf{u}, \dot{\mathbf{u}}), \quad (2.14)$$

where a labels neighbouring periodic members and $\varphi \in S^1$ locates the phase along one member. At $a = 0$, the tangent of W_j agrees with the selected linear modal plane. At finite amplitude, the embedding bends so that the full nonlinear vector field is tangent to it. This picture supplies the intuition; the computed object used later is the corrected periodic branch, and any manifold statement remains conditional on its independent return Jacobian having the required rank.

2.2.3 Finite-amplitude modes change their choreography

The double-pendulum families in Figure 2.5 make the correction visible. Near the downward equilibrium, the shape-space cycles lie close to the two straight linear modal directions. At larger energy they bend away from those lines, and their phase projections cease to be ellipses. The links remain coherently organized as an in-phase or anti-phase family, but coherence no longer means proportional sinusoids.

This is the useful meaning of *unison* here: the motion has one independent oscillatory clock, so all physical coordinates are functions of one phase along the selected family. It does not require fixed amplitude ratios, simultaneous turning points, exact zero or π phase difference at finite amplitude, or one sinusoid copied into every coordinate.

A backbone curve summarizes the same family by plotting its frequency

$$\Omega_j = \frac{2\pi}{T_j} \quad (2.15)$$

against an orbit label such as energy, action, or amplitude. The linear eigenfrequency is only the zero-amplitude limit, $\Omega_j \rightarrow \omega_j$. Moving along the backbone can change the period and the complete waveform, including the harmonics and phase relation between coordinates. A backbone organizes the computed modal family; it neither classifies every free motion nor establishes stability by itself.




Figure 2.5: Representative corrected periodic members of the in-phase (top) and anti-phase (bottom) double-pendulum families. The left panels show the joint choreography in shape space; the right panels show a phase projection. The smallest cycles approach the linear modal seeds, whereas finite-amplitude cycles deform and become non-elliptic. These selected members do not exhaust the energy surface and do not establish a global invariant manifold.

2.2.4 Action and phase only after the modal family is chosen

The one-degree action-angle picture can be recovered locally on a selected two-dimensional conservative modal surface. Intersecting that surface with a regular energy level gives a closed periodic member. An action labels nearby members and an angle clocks motion along each one, so locally

$$\dot{J}_j = 0, \quad \dot{\varphi}_j = \Omega_j(J_j). \quad (2.16)$$

This statement does not assert global integrability of the full multi-degree system. The nonlinear mode first identifies where the coherent oscillation lives; action and phase then describe the regular cycles on that selected object.

The distinction prepares the locomotor construction. The carrier-closed 3SEG model has two effective transverse coordinates and therefore needs modal organization analogous to the in-phase (IP) and anti-phase (AP) double-pendulum families. The analogy is between oscillatory dimensions: the double pendulum has two configuration degrees of freedom, whereas the complete 3SEG has five configuration coordinates and four admissible velocity coordinates, of which two form the effective transverse oscillator after carrier elimination. Families continued from the IP and AP seeds organize candidate internal supports. They do not by themselves establish return of the full opened locomotor state or a nonidentity pose increment. The next step is therefore to separate the periodic internal state from the pose that locomotion must leave open.

Table 2.1: Face-to-face structural analogies. The comparison concerns the oscillatory support exposed inside each constrained locomotor, not equality of mechanisms or of their full numbers of degrees of freedom.

Question	Classical oscillator	Locomotor counterpart
One transverse coordinate	Simple pendulum: one phase-plane loop is selected by a regular energy level; action labels neighbouring loops.	2SEG after carrier elimination: one effective transverse coordinate and rate are organized by $E_{\perp}$; reopening the propulsive channel adds the missing exchange–return question.
Several transverse coordinates	Double pendulum: total energy leaves a three-dimensional level set; in-phase and anti-phase modal families select coherent oscillations.	3SEG after carrier elimination: two effective transverse coordinates require full vector return and modal organization continued from in-phase and anti-phase seeds.
What the analogy does not provide	Neither energy nor modal identity says anything about body transport.	The opened locomotor must still solve its full reduced field, return exactly, and reconstruct a nonidentity pose increment.

2.3 The oscillator–locomotor analogies

The pendulum examples are not detached tutorials. They reveal the two dimensional transitions that reappear in the locomotors. Table 2.1 places the objects face to face. The two-segment locomotor (2SEG) is analogous to the simple pendulum only after carrier elimination has exposed one effective transverse oscillatory degree of freedom. The complete locomotor still has four configuration coordinates and three admissible velocity coordinates. Likewise, the three-segment locomotor (3SEG) is not literally a double pendulum: the analogy concerns its two effective transverse oscillatory degrees of freedom and the consequent need for modal selection.

Constraints written as $A(q)\dot{q} = 0$ allow precisely the velocity directions in the null space, or kernel, $\ker A(q)$. Three dimensions must therefore remain distinct:

$$\begin{aligned}
 \text{configuration dimension} & : \text{pose and joint coordinates,} \\
 \text{admissible-velocity dimension} & : \dim \ker A(q), \\
 \text{effective transverse dimension} & : \text{oscillatory coordinates left after carrier elimination.}
 \end{aligned} \tag{2.17}$$

For 2SEG these dimensions are 4, 3, and 1; for 3SEG they are 5, 4, and 2. This is why it is accurate to call the first construction scalar without calling the complete two-segment locomotor a one-degree-of-freedom system.

2.4 Reduced periodicity leaves the pose open

An articulated locomotor can repeat its internal motion without returning to the same place. There is no contradiction: internal deformation and placement in the world are different variables, and they obey different closure conditions. The distinction is most transparent when the configuration is written as

$$q = (g, r), \quad g \in G, \quad r \in \mathcal{R}. \tag{2.18}$$

Here g is the body pose and r is the internal shape. For planar locomotion, $G = \text{SE}(2)$: a pose contains two translations and one rotation. The shape contains joint angles or other coordinates measured relative to the body. Changing g moves the whole mechanism; changing r deforms it.

Figure 2.6 shows the split before any reduction is performed. The two body drawings can have the same internal shape while occupying different poses; a return condition on shape alone cannot distinguish them.




Figure 2.6: The configuration separates into a body pose $g \in G$ and an internal shape $r \in \mathcal{R}$. Pose describes placement in the world; shape describes deformation relative to the body. A locomotor cycle may close in its internal variables while remaining open in pose.

This split is useful when the mechanical laws are unchanged by a common rigid displacement of the mechanism. Absolute position and heading can then be removed from the equations that determine the internal motion. They are recovered afterwards by reconstruction, a standard separation in geometric and nonholonomic locomotion (Kelly and Murray, 1995; Ostrowski and Burdick, 1998; Hatton and Choset, 2011; Bloch, 2015).

2.4.1 Reduced dynamics and pose reconstruction

Velocity on the pose group is expressed in a body-attached frame. If ξ denotes this body velocity, then

$$\widehat{\xi} = g^{-1}\dot{g}, \quad \dot{g} = g\widehat{\xi}. \quad (2.19)$$

The first equality converts the world motion into a body velocity; the second reconstructs the pose from that velocity. On $\text{SE}(2)$, the components of ξ are the two body-frame translation rates and the yaw rate. Because rotations and translations do not commute, the final pose depends on the ordered history of these velocities, not merely on their separate scalar integrals.

The word *reconstruction* is reserved for this group integration. It is not used later for carrier elimination, for a closed-to-open lift, or for restoration of the propulsive channel. Those are operations on the mechanical field; pose reconstruction begins only after the reduced trajectory is known.

The reduced mechanical state is denoted by z . It contains every variable needed to advance the pose-independent dynamics: shape, shape rates and, when required, admissible carrier velocities or reduced momenta. In the unforced conservative setting, reduction and reconstruction have the schematic form

$$\dot{z} = f_0(z), \quad \xi = \Xi(z), \quad \dot{g} = g\widehat{\Xi(z)}. \quad (2.20)$$

The first equation decides which timed reduced trajectories are mechanically possible. The last two decide where one such trajectory carries the body. Reconstruction is therefore kinematic bookkeeping driven by a mechanical solution; it is not an additional force law.

In a purely kinematic model, the middle relation may reduce to a local connection, $\xi = -A(r)\dot{r}$. That compact formula is not universal. In an inertial or nonholonomic locomotor, Ξ can also depend on carrier velocity, momentum or other components of z . This is why the returned object used in this thesis is a loop in reduced *state* space rather than merely a curve in shape space.

2.4.2 Exact reduced return and nontrivial group displacement

Let a solution of Equation (2.20) start from (g_0, z_0) . It is a transporting relative-periodic motion if, for some $T > 0$,

$$z(T) = z_0, \quad \Delta g := g_0^{-1}g(T) \neq e_G, \quad (2.21)$$

where e_G is the identity pose. The first equality is exact reduced return: after one period the mechanism has the same internal configuration and the same reduced velocities. The second condition states that reconstruction has accumulated a nonidentity pose increment. This is relative periodicity: the motion is periodic after quotienting the pose symmetry, but not in the full configuration (Fasso et al., 2020).

The complete endpoint diagram was introduced in Figure 1.2. A point on its reduced loop represents the same mechanical state at the beginning and end of the period, while the two body drawings differ by the reconstructed group increment.

Full-state periodicity would instead impose both $z(T) = z_0$ and $g(T) = g_0$. It would therefore force $\Delta g = e_G$: the mechanism could move during the cycle, but it would return to its initial pose and produce no net transport per cycle. Ordinary full-state periodicity is consequently the wrong endpoint condition for a transporting gait. Pose must be removed from the return residual and evaluated as an output of reconstruction.

When the same reduced cycle is repeated, symmetry makes the body-frame increment repeat as well. The poses at complete cycles satisfy

$$g(nT) = g_0(\Delta g)^n. \quad (2.22)$$

A translational increment produces repeated advance; an increment containing rotation composes into a turning motion. The group product, rather than vector addition of world displacements, preserves the order in which translations and rotations are accumulated.

2.4.3 How planar pose increments are compared

Only a small part of the $SE(2)$ algebra is needed later. Write a planar pose as a rotation and a world position, $g = (R, p)$. Composition and inversion are

$$(R_1, p_1)(R_2, p_2) = (R_1R_2, p_1 + R_1p_2), \quad (R, p)^{-1} = (R^T, -R^Tp). \quad (2.23)$$

The rotation of the first pose therefore acts on the translation of the second. This is why two world-coordinate displacement vectors cannot simply be subtracted when their headings differ.

A constant body velocity $\xi = (v_x, v_y, \omega)$ acting for time T produces the group increment

$$\Delta g = \text{Exp}\left(T\hat{\xi}\right). \quad (2.24)$$

Conversely, the principal logarithm $\text{Log}(\Delta g)$ expresses an increment as one local twist vector. Its rotational component uses the principal angle, so the branch convention must be declared when a rotation approaches $\pm\pi$. The increments used below remain on that declared branch.

For example, comparing a cycle with a reference motion means composing group elements,

$$\Delta g_{\text{inc}} = \Delta g_{\text{base}}^{-1} \Delta g_{\text{cycle}}, \quad (2.25)$$

and then, if a vector summary is useful, taking $\text{Log}(\Delta g_{\text{inc}})$. This produces a reference-dependent residual transform. It is neither an additive energy partition nor a causal attribution of propulsion.

The equality $z(T) = z_0$ is stronger than $r(T) = r(0)$. Returning the joint angles while reversing a joint rate lands at the opposite crossing of the same oscillation. Returning the shape while changing a carrier velocity does not reproduce the same mechanical event either. A repeatable locomotor motion therefore requires a *timed* reduced-state trajectory, including the velocities needed by the model, and not only the geometric trace of the shape.

2.4.4 A worked geometric loop: displacement without selection

A two-shape toy example makes the conditional nature of reconstruction explicit. Let $x \in \mathbb{R}$ be an abelian body displacement and $r = (r_1, r_2)$ two shape coordinates. Consider the local reconstruction law

$$\dot{x} = -A_1(r)\dot{r}_1 - A_2(r)\dot{r}_2, \quad A_1 = -\frac{r_2}{2}, \quad A_2 = \frac{r_1}{2}. \quad (2.26)$$

For a closed, positively oriented loop γ bounding a region S , integration gives

$$\begin{aligned} \Delta x &= - \oint_{\gamma} (A_1 dr_1 + A_2 dr_2) \\ &= - \iint_S \left(\frac{\partial A_2}{\partial r_1} - \frac{\partial A_1}{\partial r_2} \right) dr_1 dr_2 = - \text{Area}(S). \end{aligned} \quad (2.27)$$

A loop can therefore close exactly in shape and still translate the body. Reversing its orientation reverses the displacement. This is the smallest worked example of geometric phase: a prescribed internal loop and a reconstruction law determine a pose output.

The example deliberately proves no naturality statement. It specifies no inertia, potential, or passive vector field for $r(t)$; the same geometric loop can be traversed with arbitrary timing. Moreover, the group $\mathbb{R}$ is commutative. For the finite $\text{SE}(2)$ motions below, rotation and translation must be composed in time and the exact group equation is integrated instead of being replaced by an area formula. The toy loop solves the displacement problem, not the mechanical-selection problem.

2.4.5 Reconstruction does not certify dynamical compatibility

Given a prescribed loop $r(t)$ and a reconstruction law, geometric locomotion can compute a pose increment. Many different loops can yield nonidentity increments, and they can be drawn by hand, optimized for distance, imposed by actuators or produced by feedback. This makes reconstruction a powerful tool for gait analysis and design. It does not show that a loop would occur as an unforced motion of the mechanism.

Figure 2.7 keeps that logical direction visible: the loop is supplied on the left, and pose is computed on the right. No arrow in the diagram asserts that the unforced dynamics selects the loop.




Figure 2.7: Conditional reconstruction: a closed shape loop can produce an open pose path and therefore a net pose displacement. The loop is supplied and the pose path is computed; dynamical compatibility with the unforced field must be tested separately.

Timing is part of this distinction. The same geometric curve traversed with a different time law has different shape velocities and accelerations. In an inertial system it generally requires different forces and need not satisfy the same equations. Even a closed curve in $(r, \dot{r})$ is only a candidate until its tangent agrees with the reduced vector field at every point.




Figure 2.8: Common projected Lagrange–d’Alembert construction. Both locomotors parameterize admissible velocities by a constraint–kernel basis, project away ideal reactions, integrate the reduced field, and reconstruct pose. Carrier elimination used later is a further algebraic construction inside this reduced mechanics, not a second meaning of reconstruction.

The conservative locomotor test therefore has three logically separate rows:

$$\underbrace{\dot{z}(t) = f_0(z(t))}_{\text{dynamical compatibility}}, \quad \underbrace{z(T) = z(0)}_{\text{exact reduced return}}, \quad \underbrace{\Delta g \neq e_G}_{\text{nontrivial transport}}. \quad (2.28)$$

The first row distinguishes a natural conservative motion from an arbitrarily prescribed choreography. The second makes the motion repeatable in the reduced mechanics. The third distinguishes locomotion from oscillation in place. Family or modal organization then relates neighboring solutions; it does not replace any of these three tests.

The next section turns reduced return into a computable residual. Its phase and continuation conditions choose numerical representatives of cycles and families; they do not add physical constraints to Equation (2.28).

2.4.6 The common constrained-mechanics construction

The two locomotors differ in dimension, but their equations are obtained by the same operation. Let smooth, bilateral velocity constraints be written

$$A(q)\dot{q} = 0. \quad (2.29)$$

On a constant-rank domain, choose a smooth matrix $H(q)$ whose columns span the admissible distribution,

$$\dot{q} = H(q)\eta, \quad \text{Im } H(q) = \ker A(q). \quad (2.30)$$

The Lagrange–d’Alembert balance contains the unknown ideal reaction $A(q)^\top \lambda$:

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = Q + A(q)^\top \lambda. \quad (2.31)$$

Multiplication by $H(q)^\top$ eliminates that reaction because $A(q)H(q) = 0$:

$$H(q)^\top \left(\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q \right) = 0. \quad (2.32)$$

The projected equations determine the admissible velocity coordinates η . After their pose-independent part has been integrated, the pose is reconstructed from the corresponding group-velocity components.

Figure 2.8 shows this one workflow and its two instances. The matrix is called H for 2SEG and N for 3SEG, but both are kernel bases with the same mechanical role.

The zero-work statement follows from the same geometry:

$$\dot{q}^\top A(q)^\top \lambda = (A(q)\dot{q})^\top \lambda = 0. \quad (2.33)$$

Ideal reactions therefore supply no total energy. They nevertheless define the admissible distribution, and together with the kinetic metric they structure the coupling through which energy can be redistributed between admissible channels. The full derivations are given for 2SEG in [Chapter B](#) and for 3SEG in [Chapter D](#); the shared projection argument and its scope are collected in [Chapter A](#).

The scope is important once, here. The construction assumes smooth, continuously active, bilateral, ideal, constant-rank constraints. Impacts, unilateral contact, mode switching, stick–slip, and fluid memory require different mechanics. The common method is general across the two continuous locomotors studied in this thesis, not across all possible locomotion systems.

2.5 From a returned cycle to a corrected family

A reduced periodic orbit is a closed trajectory, not a distinguished point. Choosing its initial condition only chooses where the clock reads zero. A numerical method must remove this arbitrary phase before it can distinguish a failure to close from a harmless shift along the same orbit. It must then follow neighboring corrected cycles even when a familiar physical parameter ceases to be a good coordinate. Sections, shooting and pseudo-arclength continuation address these two geometric difficulties.

2.5.1 An oriented section turns a cycle into a fixed point

Consider a reduced autonomous system

$$\dot{z} = f(z; \lambda), \quad (2.34)$$

where λ is a model or family parameter and $\Phi_\lambda^t(z_0)$ denotes the flow. Rather than observing a periodic orbit at every instant, choose a codimension-one surface $h(z) = 0$ that the flow crosses transversely. Keeping one direction of crossing gives the oriented section

$$\Sigma^+ = \{z : h(z) = 0, Dh(z)f(z; \lambda) > 0\}. \quad (2.35)$$

The strict sign has two roles. It excludes tangency, where the return event is poorly conditioned, and it discards the crossing in the opposite direction. For an oscillator, a section such as $r = 0$ is usually crossed once with $\dot{r} > 0$ and once with $\dot{r} < 0$. Retaining both can confuse a half-cycle with a full return.

Let $\tau(z)$ be the first positive return time to Σ^+ . The return map is

$$P_\lambda(z) = \Phi_\lambda^{\tau(z)}(z), \quad z \in \Sigma^+. \quad (2.36)$$

A cycle becomes the fixed-point equation $P_\lambda(z_*) = z_*$. The section does not create periodicity; it selects one representative event on the cycle and expresses all remaining closure information transversely ([Kuznetsov, 2004](#)).

[Figure 2.9](#) shows why orientation and transversality are both needed. The trajectory meets the geometric section twice, but only one crossing represents the chosen complete-cycle event.

For a relative-periodic locomotor, the map acts on the reduced state. The pose is reconstructed over the accepted return segment and is not included in the fixed-point equality. Including absolute pose in the return vector would replace relative periodicity by full-state periodicity and force the net pose increment to the identity.

2.5.2 Shooting corrects closure and a phase row fixes the clock

A section is convenient when a clean event is available. A more general description treats the initial state z_0 and the period T as unknowns. Single shooting integrates once around the proposed cycle




Figure 2.9: An oriented Poincaré section samples one crossing per complete cycle. The continuous orbit becomes a fixed point of the return map; the crossing with the opposite orientation is discarded.

and forms

$$R(z_0, T, \lambda) = \Phi_\lambda^T(z_0) - z_0. \quad (2.37)$$

The equation $R = 0$ closes the trajectory, but it has no preferred time origin: if $z_*(t)$ is periodic, every shifted point $z_*(t + s)$ represents the same orbit. This neutral tangent direction makes a Newton correction underdetermined unless one scalar phase condition is added. A common local choice is

$$\psi(z_0) = \langle z_0 - z_{\text{ref}}, f(z_{\text{ref}}; \lambda) \rangle = 0. \quad (2.38)$$

It asks the new initial point to lie in a hyperplane transverse to a nearby reference orbit. Other pointwise or integral phase conditions serve the same purpose. The corrected cycle solves

$$\mathcal{F}(z_0, T, \lambda) = \begin{bmatrix} R(z_0, T, \lambda) \\ \psi(z_0) \end{bmatrix} = 0. \quad (2.39)$$

The phase row is a gauge: it chooses the clock origin but exerts no force and does not imply stability. An oriented section and a phase condition are two representations of the same need. The former chooses the sampling event geometrically; the latter pins a representative inside a shooting or boundary-value unknown.

Correction is local. A linear mode, a small-amplitude orbit or a neighboring corrected solution can provide a seed close to a root. Newton iterations then adjust the initial state and period until the scaled closure and phase residuals meet their declared tolerances. The seed supplies a tangent or an initial guess; it is not the finite-amplitude solution. This distinction is especially important for nonlinear modal families, whose waveform and period can depart substantially from the linear seed (Peeters et al., 2009).

2.5.3 Pseudo-arclength follows the curve through a fold

After one corrected orbit has been obtained, neighboring orbits are represented as roots of a residual

$$F(u, \lambda) = 0, \quad (2.40)$$

where u collects the orbit unknowns, including its period and phase-fixed representation. Stepping λ and correcting u works while the branch is a graph $u(\lambda)$. At a fold, the solution curve remains smooth in the full (u, λ) space but turns in its projection onto λ . The physical parameter has failed as a local branch coordinate; the family has not necessarily ended.

Pseudo-arclength continuation follows the curve itself. With $y = (u, \lambda)$, let t_j be a unit tangent at a known root y_j . A predictor and its corrected point satisfy

$$y_p = y_j + \Delta s t_j, \quad F(y) = 0, \quad t_j^\top (y - y_p) = 0. \quad (2.41)$$

The first expression advances along the tangent. The last equation restricts the correction to the hyperplane normal to that tangent. It replaces the failed choice of λ by a local branch gauge and can pass a simple fold. A branch point is different: several local branches meet and additional information is required to decide whether and how to switch branches.

Figure 2.10 depicts this geometry. The predictor advances along the solution tangent, and the augmented correction returns to the residual zero set without requiring the plotted parameter to remain monotone.




Figure 2.10: Pseudo-arclength predictor–corrector geometry. A tangent predicts a nearby point on the solution curve; correction then solves the physical residual together with a hyperplane gauge. The gauge enables continuation through a fold in the plotted parameter but adds no mechanical equation.

Continuation therefore follows the residual it is given. It can organize a family, expose folds and supply neighboring initial guesses, but it cannot decide by itself whether the solutions are natural, transporting, modal or stable. Those meanings depend on the vector field, the return rows, reconstruction and the claim-specific tests attached to each solution.

2.5.4 When a computed family is locally a manifold

The word *manifold* concerns the solution set, not the smooth appearance of one plot. After phase fixing, suppose the independent orbit equations take the form

$$F : \mathbb{R}^{N+1} \rightarrow \mathbb{R}^N, \quad F(y) = 0. \quad (2.42)$$

If F is smooth and

$$\text{rank } DF(y_\star) = N, \quad (2.43)$$




Figure 2.11: Local family geometry. Left: a full-rank phase-fixed return residual defines a local solution curve. Centre: the observable or chosen continuation parameter can fold while that curve remains regular. Right: a branch point can destroy the single-chart description. A point cloud or an observable image alone establishes none of these rank properties.

then the implicit-function theorem gives a smooth one-dimensional zero set near y_* . This is the phase-fixed branch. Restoring the arbitrary phase $\varphi \in S^1$ along each nontrivial cycle produces locally an object of the form

$$I \times S^1, \quad (2.44)$$

provided neighbouring cycles remain embedded, have a regular phase action, and belong to one fixed autonomous vector field. The interval I labels the cycle; the circle labels phase on that cycle.

Figure 2.11 distinguishes three geometries. A regular phase-fixed branch supplies a local chart. A fold in one observable can occur while the branch remains regular in its full solution space. At a branch point or other rank loss, several tangents can meet and a single smooth chart need not continue through the event.

An observable map O introduces a second distinction. Even if the underlying solution family is smooth, O can fold it, make distant points intersect in the plot, or lose rank. The plotted image is therefore called an observable projection or chart, not the manifold by definition. The NLM and RLM language used later refers to regular local solution sets in their declared orbit spaces; no global union across branches, modal sectors, or models is inferred from a visualization.

2.5.5 Candidate, corrected cycle, and mechanically qualified member

Three statuses keep the numerical construction separate from its mechanical interpretation.

Candidate. A seed, predictor or approximate cycle lies near the expected branch. Exact closure, dynamical accuracy and claim-specific meaning have not yet been established.

Corrected cycle. The declared orbit residual and phase row pass with stated scaling and tolerances. This establishes a numerical root of those equations, but not stability, uniqueness, global continuation, or the mechanical interpretation attached to the branch.

Mechanically qualified member. The corrected cycle also satisfies the complete field, parameter identity, reduced-return, reconstruction, and energy, exchange, or modal conditions required by the stated result. This status is specific to the claim: it does not turn a numerical construction into a theorem, an exhaustive classification, a robustness result, or an experiment.

A smooth branch plot can contain candidates that have not passed the return test, and a small shooting residual can omit a necessary mechanical row. Conversely, a finite collection of qualified members remains a numerical construction rather than a proof of a global manifold. The conclusion attached to a result must therefore follow the equations and checks actually used, not the visual smoothness or size of the plotted family.

The conservative chapters use this hierarchy in a fixed order: obtain a seed, correct exact reduced return, continue neighboring solutions, reconstruct their pose increments and finally apply the dimensional exchange or modal tests appropriate to the model. The numerical gauges make the

family computable; the mechanics gives it meaning.

2.6 What losses change

The preceding constructions concern an unforced conservative field. If a state is initialized exactly on one of its periodic or relative-periodic solutions, the ideal model contains no loss that must be replenished. The orbit is invariant under exact initialization; it is not thereby attractive, stable, or robust to a change of model.

2.6.1 From energy conservation to work balance

A damped, forced nonlinear oscillator provides the smallest worked example:

$$m\ddot{x} + c\dot{x} + kx + \alpha x^3 = F \cos(\Omega t). \quad (2.45)$$

With

$$E = \frac{1}{2}m\dot{x}^2 + \frac{1}{2}kx^2 + \frac{1}{4}\alpha x^4, \quad (2.46)$$

multiplication by $\dot{x}$ gives

$$\dot{E} = \underbrace{F \cos(\Omega t)\dot{x}}_{P_{\text{in}}} - \underbrace{c\dot{x}^2}_{P_{\text{diss}}}. \quad (2.47)$$

Without input, $P_{\text{in}} = 0$ and a non-equilibrium motion loses energy. With periodic input, stored energy may rise and fall within a cycle. If the state returns after a period T , however, its net change is zero and

$$z(T) = z(0) \implies \int_0^T P_{\text{in}} dt = \int_0^T P_{\text{diss}} dt. \quad (2.48)$$

This is an exact necessary accounting identity for a periodic response. It does not determine the waveform, phase, stability, or existence of that response; the complete forced equations and return conditions still do so.

Figure 2.12 contrasts the three energy regimes. The sustained trace remains periodic because input work replaces losses, not because damping has somehow become conservative.

2.6.2 The forcing clock defines the return

Let the forcing period be $T_{\text{exc}} = 2\pi/\Omega$. Observing the system at one fixed phase of the input defines a stroboscopic map

$$P_{A,T} : z_0 \mapsto \Phi_{A,T}^T(z_0). \quad (2.49)$$

A synchronized periodic response is a fixed point,

$$G(z_0, A, T) := P_{A,T}(z_0) - z_0 = 0. \quad (2.50)$$

Unlike an autonomous orbit, this response is referred to an imposed forcing phase. The input clock removes the continuous time-shift freedom in the stroboscopic representation. An autonomous forcing lift can restore an explicit phase state, but it represents the same period-matched response.

The forced return equation, not an amplitude or phase plot, defines the response. If $z_0 \in \mathbb{R}^n$ and the derivative of G has rank n , then allowing amplitude A and period T to vary gives locally a two-dimensional synchronized-response sheet $\mathcal{S}_R$ in orbit-forcing space. Pose reconstruction selects $\mathcal{M}_R = \{s \in \mathcal{S}_R : \Delta g(s) \neq e_G\}$. Regular local components of this transported subset, continued within the one-to-one resonance construction, are called the Resonant Locomotion Manifold (RLM). Here *resonant* names the family organized around the one-to-one resonance; it does not make exact fundamental quadrature a definition. Mean speed, waveform, phase, and reconstructed pose are observable maps; their plotted image may fold or self-intersect even when the underlying solution set is regular.




Figure 2.12: Changing the energy regime changes the repeatable object. A conservative oscillation can persist at fixed energy, a damped free motion decays, and a maintained periodic response requires input work. The last case is forced rather than passive, even when the forcing accompanies the response instead of prescribing its full waveform.

2.6.3 Why quadrature is useful but not sufficient

In the linear limit of Equation (2.45), write

$$x(t) = X \cos(\Omega t - \phi), \quad (2.51)$$

where ϕ is the displacement lag relative to the force. Near the familiar lightly damped resonance, $\phi \approx \pi/2$. Then $\dot{x} \approx \Omega X \cos(\Omega t)$: the force is in phase with velocity, and the mean fundamental input power is positive. With the output-minus-input convention used later,

$$\phi_{x-F} = \arg \hat{x}_1 - \arg \hat{F}_1 \approx -\frac{\pi}{2}. \quad (2.52)$$

Reversing the order of signals reverses the sign; the mechanical content is the quarter-period relation.

At finite amplitude, nonlinear inertia, stiffness, constraints, and several loss channels can bend the response branch and generate harmonics. The forcing frequency is no longer compared with one immutable natural frequency, and the response need not be a graph over frequency. Quadrature of the fundamental is therefore a signal-dependent local marker. It can reveal that the fundamental torque supplies work in the right timing, but it is neither pointwise power balance nor a universal optimum. Chapter 6 consequently keeps two statements separate: the cycle-integrated work balance is exact for every periodic response, whereas the observed alignment of quadrature and a speed ridge is local to the reported part of the response family.

The same distinction guides feedback. Accompanying a response family does not require prescribing every harmonic when a few observables locate the desired portion of the solution sheet. In the controller studied here, deformation–torque phase adjusts forcing frequency on a fast loop, and mean forward speed adjusts forcing amplitude on a slower loop. This is local phase–speed regulation for the forced model, not stabilization of the conservative NLM.

2.7 The order of the construction

The foundations now assemble into a fixed scientific order. A curve does not advance from one step to the next merely because it looks oscillatory or reconstructs a displacement.

1. **Dynamical compatibility.** The time-parameterized reduced trajectory, including velocities, solves the specified mechanical field. A linear eigenvector, guessed sinusoid, or prescribed shape loop is only a candidate.
2. **Nontrivial exact reduced return.** The corrected trajectory returns to the same reduced state under a declared phase or minimal-period convention. Straight relative equilibria provide seeds and boundaries but are not converted into cycles by an arbitrary return time.
3. **Declared locomotor output.** Pose reconstruction evaluates the group increment. A nonidentity increment establishes a moving relative-periodic motion; attributing an additional propulsive effect to internal oscillation requires comparison with a matched straight-motion reference.
4. **Family organization and modal identity.** Continuation follows connected branches of corrected solutions. When several transverse coordinates are present, cycle-level metrics distinguish organizations continued from in-phase and anti-phase seeds. These labels organize qualified cycles; they do not replace return.
5. **Dissipative maintenance and regulation.** When the field includes losses and input, a stroboscopic return and cycle work balance define the maintained response. Phase and mean speed can then serve as local feedback coordinates without redefining conservative naturality.

The closed–open exchange construction enters at a precise location inside the second step. In the scalar two-segment problem, once the other return rows are imposed, integrated exchange of a derived transverse storage can replace and interpret the last return row under the theorem’s section hypotheses. In the three-segment problem, one storage value cannot identify a point on the higher-dimensional section. Full vector return remains necessary; exchange then becomes a balance identity and diagnostic, while modal metrics classify the returned cycle.

This order prevents three shortcuts. Pose reconstruction is not a test of the mechanical field. Zero exchange is not an independent equation added after complete return. Modal identity is not a substitute for return. With these roles fixed, [Chapter 3](#) can follow the first two-segment family from a straight relative-equilibrium seed to a corrected finite-amplitude cycle and its reconstructed pose.

Bridge to the conservative construction

The chapter has assembled a single causal sequence. Energy and action organize regular one-degree oscillations; nonlinear modes recover a corresponding local organization after a modal family has been selected in several degrees of freedom. Locomotion then replaces full-state periodicity by exact reduced return and treats pose as a reconstructed output. Shooting fixes the return, continuation follows neighbouring roots, and local rank explains when those roots form a smooth branch or a phase-swept local manifold chart. Constraint projection supplies the common mechanical field to which these tools are applied.

[Chapter 3](#) now instantiates this sequence on the two-segment locomotor. [Chapter 4](#) will expose the effective transverse oscillator and propulsive channel that make the pendulum analogy mechanically precise; [Chapter 5](#) will show exactly where the scalar analogy stops.

Part II

Natural locomotion under total-energy conservation

Chapter 3

Natural-locomotion families of the two-segment locomotor

The simplest locomotor considered in this thesis has one internal joint. It is therefore the first setting in which the pendulum analogy can be made mechanical rather than verbal: after the carrier motion has been separated, one effective transverse degree of freedom remains. The complete locomotor is nevertheless not a one-degree-of-freedom oscillator. Its admissible velocity contains forward, yaw, and joint components, and a returned internal motion may accumulate a planar pose increment.

This chapter introduces the *two-segment locomotor (2SEG)*, derives its pose-reduced conservative field, and follows an oscillatory direction from a straight relative equilibrium into a finite-amplitude family. The resulting cycles do not share one prescribed clock or choreography: their period, waveform, and inter-coordinate phase change along the branch. Pose reconstruction then shows what each returned reduced cycle does in the plane, while a matched straight motion separates transport of the moving carrier from the smaller path modulation associated with the oscillation.

Continuation of periodic families is an established tool in nonlinear dynamics and locomotion (Remy, 2011; Raff et al., 2022; Kerschen et al., 2009; Peeters et al., 2009). The result specific to this chapter is the formulation of complete reduced return with open planar pose, followed as a locomotor branch of the declared constrained mechanics. The 379 exported members document its observable trends and reconstructed pose. Their stored endpoints do not, by themselves, reproduce the unavailable raw collocation residual or a branch-wide rank test; Chapter 4 independently re-integrates the field and evaluates local rank at 29 matched roots. Neither evidence layer establishes global existence, attraction, stability, or one global smooth manifold. The straight motions used as seeds are relative equilibria, not periodic cycles created by assigning them an arbitrary period.

3.1 Mechanical model

3.1.1 Configuration, constraint, and pose reconstruction

The 2SEG locomotor consists of a planar chassis and one internal rigid segment joined by a torsional spring. Its configuration is

$$q = (g, \delta), \quad g = (x, y, \theta) \in \text{SE}(2), \quad (3.1)$$

where δ is the relative joint angle. The chassis can roll along its heading but cannot move sideways at its contact point. Figure 3.1 fixes the positive directions and reduced variables before the constraint is written.

$$A(\theta)\dot{q} = 0, \quad A(\theta) = \begin{bmatrix} -\sin \theta & \cos \theta & 0 & 0 \end{bmatrix}. \quad (3.2)$$




Figure 3.1: The two-segment locomotor and its reduced variables. The ideal wheel-like constraint removes lateral contact velocity; it does not remove forward speed or yaw. The spring-mounted internal segment is described by joint angle δ and joint rate σ .

This is called a *nonholonomic* constraint: it restricts instantaneous velocities and cannot, in general, be integrated into a fixed relation among x , y , and θ . An admissible velocity can be written as

$$\dot{q} = H(\theta)\eta, \quad H(\theta) = \begin{bmatrix} \cos \theta & 0 & 0 \\ \sin \theta & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \eta = [v \quad \omega \quad \sigma]^\top, \quad (3.3)$$

with longitudinal speed v , yaw rate $\omega = \dot{\theta}$, and joint rate $\sigma = \dot{\delta}$. Because $A(\theta)H(\theta) = 0$, this representation enforces the lateral constraint by construction (Bloch, 2015; Rajaomarasata et al., 2025a).

Task-oriented models of nonholonomic mobile manipulators likewise parameterize available velocities to expose task and redundant directions (De Luca et al., 2007). Here the same kernel parameterization serves a different, dynamic purpose: it projects the Lagrange–d’Alembert equations and eliminates the ideal reaction forces.

This is the first application of the constrained-reduction pattern used again for the three-segment locomotor. Starting from the Lagrange–d’Alembert equation with reaction $A(q)^\top \lambda$, multiplication by $H(q)^\top$ gives

$$H(q)^\top \left[\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q \right] = 0, \quad \text{Im } H(q) = \ker A(q). \quad (3.4)$$

The reaction disappears because every column of H is an admissible velocity. Equivalently, an ideal reaction has zero power on every admissible motion. It does not energize the locomotor; the constraint distribution and the kinetic metric instead determine how its admissible velocity components are coupled.

Table 3.1 makes the dimensions explicit. The same four operations—constraint, kernel basis, projected dynamics, and later pose reconstruction—will be used with the matrix N in Chapter 5.

The reusable scope is precise: the projection applies to smooth, bilateral, ideal, continuously active, constant-rank velocity constraints. It does not by itself describe impacts, unilateral or switching contacts, or stick–slip dynamics.

The distinction between exact reduced return and net pose displacement was introduced in Figure 1.2.

The dynamics are solved in the four-dimensional state

$$z = (\delta, v, \omega, \sigma)^\top. \quad (3.5)$$

The planar pose is then reconstructed from

$$\dot{x} = v \cos \theta, \quad \dot{y} = v \sin \theta, \quad \dot{\theta} = \omega. \quad (3.6)$$

Consequently, a closed trajectory in z need not close in (x, y, θ) . That distinction is the elementary mechanism by which exact reduced return and locomotion coexist.

Table 3.1: Common projected Lagrange–d’Alembert pattern specialized to the 2SEG locomotor. The effective transverse dimension is obtained only after the two carrier velocities have subsequently been eliminated; it is not the dimension of the complete locomotor.

Operation	General object	2SEG realization
configuration	$q \in Q$	$q = (x, y, \theta, \delta)$: three pose coordinates and one joint
velocity constraint	$A(q)\dot{q} = 0$	one lateral no-slip row, of rank one
admissible basis	$\dot{q} = H(q)\eta$, $\text{Im } H = \ker A$	$\eta = (v, \omega, \sigma)$, three admissible velocities
projected mechanics	project the Lagrange–d’Alembert balance with $H^\top$	three reduced acceleration balances
transverse oscillator	carrier velocities eliminated from the kinetic energy	one effective transverse degree of freedom (δ, σ)
pose output	integrate the group equation after the reduced solve	reconstruct (x, y, θ) from (v, ω)

3.1.2 Compact reduced dynamics

The computations use the dimensionless parameterization

$$\alpha = \varepsilon, \quad \beta = \mu = 1 - \varepsilon, \quad j_1 = \gamma\alpha^3, \quad j_2 = \gamma\beta^3, \quad (3.7)$$

and the passive joint potential

$$U(\delta) = \frac{1}{2}k_2\delta^2 + \frac{1}{4}k_4\delta^4. \quad (3.8)$$

The reference values are

$$\varepsilon = 0.681, \quad \gamma = 2.504, \quad k_2 = 1, \quad k_4 = 10. \quad (3.9)$$

Applying Equation (3.4) to the kinetic and potential energies of this locomotor gives the compact balance

$$M(\delta)\dot{\eta} + c(\delta, \eta) + [0 \quad 0 \quad U'(\delta)]^\top = 0. \quad (3.10)$$

Here c contains the velocity-quadratic inertial terms and

$$M(\delta) = \begin{bmatrix} 1 & -\rho & -\rho \\ -\rho & B_{11} & B_{12} \\ -\rho & B_{12} & B_{22} \end{bmatrix}, \quad (3.11)$$

with

$$\rho(\delta) = \mu\beta \sin \delta, \quad (3.12a)$$

$$B_{11}(\delta) = j_1 + j_2 + \alpha^2 + \mu\beta^2 + 2\mu\alpha\beta \cos \delta, \quad (3.12b)$$

$$B_{12}(\delta) = j_2 + \mu\beta^2 + \mu\alpha\beta \cos \delta, \quad (3.12c)$$

$$B_{22} = j_2 + \mu\beta^2. \quad (3.12d)$$

The dimensional derivation, the explicit inertial vector, the zero-work check, and the scaling dictionary are given in [Chapter B](#); that appendix uses the same velocity basis and points back to the present reduced equations. In physical terms, ρ couples translation to the two angular rates when the joint is deflected, while the terms proportional to $v\omega$ couple a moving chassis to its yaw and internal dynamics.

Where $M(\delta)$ is nonsingular, define

$$a(\delta, \eta) = -M(\delta)^{-1} \left(c(\delta, \eta) + \begin{bmatrix} 0 & 0 & U'(\delta) \end{bmatrix}^\top \right) = \begin{bmatrix} a_v & a_\omega & a_\sigma \end{bmatrix}^\top. \quad (3.13)$$

The reduced dynamics are then

$$\dot{z} = f(z) = \begin{bmatrix} \sigma & a_v & a_\omega & a_\sigma \end{bmatrix}^\top. \quad (3.14)$$

With no forcing or damping, the mechanical energy

$$E(z) = \frac{1}{2} \eta^\top M(\delta) \eta + U(\delta) \quad (3.15)$$

is conserved.

3.2 Straight motion and the oscillatory seed

For every constant speed $\bar{v}$, the reduced state

$$z_{\text{re}}(\bar{v}) = \begin{bmatrix} 0 & \bar{v} & 0 & 0 \end{bmatrix}^\top \quad (3.16)$$

satisfies

$$f(z_{\text{re}}(\bar{v})) = 0. \quad (3.17)$$

Indeed, the spring force, the translation–rotation coupling, and every velocity-quadratic term vanish at $\delta = \omega = \sigma = 0$. In the plane, however, this state produces

$$x(t) = x_0 + \bar{v}t \cos \theta_0, \quad y(t) = y_0 + \bar{v}t \sin \theta_0, \quad \theta(t) = \theta_0. \quad (3.18)$$

It is therefore a straight translation, or relative equilibrium, rather than a motionless configuration. The family parameter also generates the exact neutral tangent

$$Df(z_{\text{re}}(\bar{v})) \begin{bmatrix} 0 & 1 & 0 & 0 \end{bmatrix}^\top = 0. \quad (3.19)$$

Internal oscillations are sought transverse to that tangent, in $\xi_\perp = (\delta_1, \omega_1, \sigma_1)^\top$. Linearization gives

$$\dot{\xi}_\perp = J_\perp(\bar{v}) \xi_\perp. \quad (3.20)$$

At zero carrier speed, the reference design has the oscillatory eigenpair

$$\lambda_\pm(0) = \pm 3.30569i. \quad (3.21)$$

The associated eigenvector fixes the relative amplitudes and phases of joint angle, yaw rate, and joint rate. At nonzero speed the corresponding pair need not remain exactly imaginary; it initializes the nonlinear solve, which then determines the finite-amplitude cycle (Shaw and Pierre, 1993; Kerschen et al., 2009).

If $\lambda = \alpha_\lambda + i\Omega$ and $p = p_r + ip_i$ denote that pair and a transverse eigenvector at a selected $\bar{v}_0$, embed p as $\hat{p} = (p_\delta, 0, p_\omega, p_\sigma)^\top$. The initial loop is

$$z_{\text{seed}}(t) = z_{\text{re}}(\bar{v}_0) + A(\hat{p}_r \cos \Omega t - \hat{p}_i \sin \Omega t), \quad A > 0. \quad (3.22)$$

Newton correction supplies the omitted nonlinear deformation, including the variation of v around its mean.

3.3 Shooting and continuation

Let $\Phi_T(z_0)$ be the flow of Equation (3.14). A reduced cycle is a state and period satisfying

$$\Phi_T(z_0) - z_0 = 0. \quad (3.23)$$

Because shifting the initial time gives the same orbit, one phase equation is added, for example

$$(z_0 - z_{\text{ref}})^\top f(z_{\text{ref}}) = 0. \quad (3.24)$$

Here z_{ref} is a point on a previously corrected nontrivial cycle, so $f(z_{\text{ref}}) \neq 0$. A straight relative equilibrium cannot be used in this tangent condition because its vector field vanishes; the first nontrivial seed instead uses a component or modal phase convention. Shooting integrates the model over one candidate period and corrects z_0 and T until these return and phase residuals vanish.

The conservative cycles form a branch rather than isolated attracting limit cycles. The numerical formulation therefore introduces an auxiliary isotropic term $\kappa\eta$ on the left-hand side of Equation (3.10):

$$M(\delta)\dot{\eta} + c(\delta, \eta) + [0 \quad 0 \quad U'(\delta)]^\top + \kappa\eta = 0. \quad (3.25)$$

This unfolding does not enlarge the set of nonstationary periodic solutions. Indeed, its exact energy balance is

$$\dot{E} = -\kappa \eta^\top \eta, \quad 0 = E(T) - E(0) = -\kappa \int_0^T \eta^\top \eta \, dt. \quad (3.26)$$

The integral is strictly positive on every nontrivial cycle, hence periodicity forces $\kappa = 0$. The extra coordinate regularizes the continuation equations without changing the conservative periodic branch. With $X = (z_0, T, \kappa)$, a tangent predictor is corrected together with the scalar pseudo-arclength condition

$$\tau_k^\top (X - X_k) = \Delta s. \quad (3.27)$$

This permits the solver to follow folds without choosing period, energy, or mean speed as a global branch parameter (Peeters et al., 2009; Kuznetsov, 2004). At finite numerical precision, the computed values satisfy

$$|\kappa| \leq 7.8 \times 10^{-12}. \quad (3.28)$$

This formulation also fixes what is meant by a branch and by a local manifold chart. Collect the four return equations and the phase equation in $\mathcal{F}_N(z_0, T, \kappa) = 0$. There are five equations for the six unknowns (z_0, T, κ) . If their phase-fixed Jacobian has rank five at a solution, the implicit-function theorem makes the local zero set a smooth one-dimensional curve. Restoring phase along the nontrivial cycles then sweeps a locally two-dimensional invariant family of the fixed conservative field, because Equation (3.26) places every exact periodic root on $\kappa = 0$, provided the curve is embedded. On the transported subset, a regular continued component is then a local chart of the Natural Locomotion Manifold (NLM).

The stored samples vary smoothly and the unfolding coordinate remains at the level reported in Equation (3.28), but the available data do not contain branch-wide singular values for this Jacobian. The result is therefore reported unconditionally as a *computed family* or *continuation branch*; no global smoothness across folds, branch points, or other model architectures is inferred. The 29 roots studied in Chapter 4 supply independent exact-field and local-rank evidence at those sampled roots only, not for every exported record or between samples.

3.4 The continued family and its locomotor meaning

For a cycle of period T , define its mean longitudinal speed by

$$\bar{v} = \frac{1}{T} \int_0^T v(t) \, dt. \quad (3.29)$$

Mean speed is a useful observable and branch label. It is not the naturality criterion: a straight relative equilibrium and many prescribed internal loops can have the same value of $\bar{v}$. The defining requirement here is that the time-parameterized reduced trajectory solves the unforced field and returns exactly, with pose reconstructed only afterwards. The historical continuation was posed to satisfy this requirement. For its 379 exported records, the population-wide checks recover endpoint coincidence, sampled energy variation, the unfolding coordinate, and pose; exact-field re-integration and numerical rank are independently available for the 29 matched roots described above. The exported continuation window is sampled at 379 points in $-200 \leq \bar{v} \leq 200$. It is shown directly in Figure 3.2, rather than represented only by its numerical ranges. The population-wide endpoint, energy, auxiliary-parameter, and pose checks are reported in Section B.4.

The two panels expose two different nonlinear changes. Period varies along the branch, and the relative participation of the carrier and internal segment also changes: the joint keeps a finite excursion while chassis yaw becomes small at large $|\bar{v}|$. Scaling one harmonic seed therefore cannot generate the displayed family.

The timing change can be stated independently of the arbitrary point at which a cycle is cut. Write a returned coordinate as $q(t) = \sum_k \hat{q}_k \exp(2\pi i k t / T)$. For $k \geq 1$, the heading coefficient reconstructed from yaw rate is

$$\hat{\theta}_k = \frac{T}{2\pi i k} \hat{\omega}_k. \quad (3.30)$$

The first-harmonic phase difference and yaw-rate total harmonic distortion (THD) are then

$$\phi_{\delta\theta} = \arg\left(\frac{\hat{\delta}_1}{\hat{\theta}_1}\right), \quad \text{THD}_\omega = \frac{\left(\sum_{k \geq 2} |\hat{\omega}_k|^2\right)^{1/2}}{|\hat{\omega}_1|}. \quad (3.31)$$

A time shift rotates both first-harmonic coefficients by the same angle, so $\phi_{\delta\theta}$ does not depend on the chosen section point. An exact single-frequency yaw motion would instead have $\text{THD}_\omega = 0$.




Figure 3.2: Exported 379-member continuation window. (a) The corrected period changes from 0.6674 near zero mean speed to about 0.799 at the displayed ends. (b) Joint excursion remains between 1.62 and 1.99 rad while yaw-rate excursion falls from about 1.94 to 0.135. L, M, and H identify the members used in Figure 3.3. The figure describes the sampled continuation component; it is not a branch-wide regularity or stability result.




Figure 3.3: The family does not repeat one fixed sinusoidal choreography. (a) Joint angle and reconstructed heading for the L, M, and H members, each shifted to maximum joint angle and normalized separately so that timing can be compared without hiding the amplitude change already reported in Figure 3.2. (b) The continuously unwrapped first-harmonic lag changes across the family; values separated by 360 deg describe the same wrapped phase. (c) The yaw-rate THD rises from 4.30% for L to 13.18% for M and 17.30% for H. Thus both inter-coordinate phase and harmonic content are branch quantities. The conclusion is specific to this computed 2SEG family and is not a comparison with all possible controlled trajectories.

Near zero speed the joint and heading fundamentals are nearly opposite; by the high-speed representative their unwrapped lag has changed by about 85 deg. At the same time the yaw-rate waveform contains a growing higher-harmonic component. A controller that prescribed one sine wave with a fixed frequency and fixed phase relation would therefore miss two features of the unforced conservative solutions: their speed-dependent clock and their speed-dependent coordination.

The pose increment is obtained after the reduced solve:

$$\Delta\theta = \int_0^T \omega(t) dt, \quad (3.32a)$$

$$\Delta x = \int_0^T v(t) \cos \theta(t) dt, \quad (3.32b)$$

$$\Delta y = \int_0^T v(t) \sin \theta(t) dt. \quad (3.32c)$$

Thus reduced return and planar closure are distinct questions. The reported exports have coincident reduced endpoints while their reconstructed positions change; for the representative rectilinear motions, the heading increment is close to zero and successive periods extend the same macroscopic direction of travel.

The distinction is visible without reading a state-space plot. The introductory example in Figure 1.3 places one closed reduced loop beside its open reconstructed pose. The carrier heading oscillates during the cycle, so the path is curved locally, while its endpoint remains displaced after the reduced state has returned. That dimensionless trajectory visualizes the reduced-return/open-pose property for the displayed member. Its endpoint and pose checks do not independently reproduce the historical collocation residual, and they establish neither stability nor hardware motion.

To separate this moving-family statement from a stronger propulsion claim, each cycle is compared with the straight relative equilibrium having the same sampled mean speed and period. With both poses expressed from the same initial carrier frame, define

$$\xi_{\text{base}} = [\bar{v} \ 0 \ 0]^T, \quad \Delta g_{\text{base}} = \text{Exp}_{\text{SE}(2)}\left(T\hat{\xi}_{\text{base}}\right), \quad \Delta g_{\text{inc}} = \Delta g_{\text{base}}^{-1} \Delta g_{\text{cycle}}. \quad (3.33)$$

Write the principal group logarithm as

$$\text{Log}_{\text{SE}(2)}(\Delta g_{\text{inc}})^\vee = (\rho_{\text{inc}}, \phi_{\text{inc}}), \quad \rho_{\text{inc}} \in \mathbb{R}^2. \quad (3.34)$$

The two entries distinguish the residual translation and rotation after the matched composition. This is a well-defined counterfactual comparison on $\text{SE}(2)$; unlike subtracting $\bar{v}T$ from a world-coordinate path, it remains meaningful when a cycle contains rotation.

Figure 3.4 applies this group comparison to the continued two-segment family and displays both the absolute and normalized residuals.

Among the 377 rows satisfying $|\bar{v}| \geq 1$, the median translational-log ratio is 0.2603%, its 95th percentile is 7.868%, and 75.33% of the rows lie below 1%; the maximum is 10.88%. Thus, away from a nearly zero matched-straight displacement, that baseline typically dominates the computed translation, with speed-dependent exceptions. Near $\bar{v} = 0$ the normalized ratio is deliberately omitted and the absolute increment must be read instead. The result is therefore a family of oscillatory motions riding on and modulating straight translation, not evidence that internal oscillation is the main source of its total displacement.




Figure 3.4: Matched-straight pose decomposition over the 379-cycle displayed window. (a) Cycle longitudinal displacement and the matched straight displacement are visually close. (b) The translation of the principal SE(2) logarithm ρ_{inc} resolves their group residual. (c) The normalized residual is reported only for $|\bar{v}| \geq 1$, because a ratio to a vanishing baseline has no useful interpretation. The residual is a reference-dependent modulation, not a unique causal or energetic measure of propulsion.

Chapter 4 independently realizes its exchange–return theorem on 29 re-integrated cycles matched to this family.

3.5 The scalar return coordinate

The full return equation closes four reduced coordinates. Its scalar structure becomes visible on the oriented section

$$\Sigma^+ = \{z : \delta = 0, \sigma > 0\}. \quad (3.35)$$

Let

$$z_0 = (0, v_0, \omega_0, \sigma_0), \quad z^+ = (0, v^+, \omega^+, \sigma^+) \quad (3.36)$$

be consecutive crossings of Σ^+ . Once the event and the two carrier coordinates satisfy

$$v^+ = v_0, \quad \omega^+ = \omega_0, \quad (3.37)$$

the remaining return equation is

$$\sigma^+ - \sigma_0 = 0. \quad (3.38)$$

The sign in the section distinguishes the two zero crossings of δ in one oscillation. Chapter 4 shows when this crossing-speed equation is equivalent to zero integrated mechanical exchange, rather than an additional condition imposed on an already closed cycle.

A transported conservative family with one scalar return left

Straight translations become equilibria only after planar pose has been removed. A transverse oscillatory eigenpair of those reduced equilibria seeds a finite-amplitude branch whose period, carrier participation, inter-coordinate phase, and harmonic content vary with the member. Nonlinear correction and continuation are therefore recovering mechanical choreographies that cannot be replaced by one rescaled sinusoid.

The returned reduced states reconstruct to open poses. Relative to a matched straight motion, their total translation is usually dominated by baseline carrier translation outside the near-zero-speed region, while the oscillation produces a smaller, speed-dependent path modulation. Natural locomotion is not being identified with large displacement or with mean speed. The unresolved mechanical point is instead the last scalar return coordinate on the positive joint-crossing section. [Chapter 4](#) shows that this coordinate has an exact energy-exchange interpretation.

Chapter 4

Effective oscillator, propulsive channel, and exact scalar return

The family in [Chapter 3](#) was computed by imposing reduced return directly. This chapter asks what the last scalar return condition means mechanically. The answer comes from separating two energy reservoirs in the same 2SEG locomotor: an *effective transverse oscillator* and the complementary *propulsive carrier channel*, shortened below to *propulsive channel*. The channel is propulsive because it contains the complementary carrier motion that contributes to pose transport; the name does not imply that internal oscillation causes all of the displacement.

The construction proceeds through four distinct operations. Carrier elimination completes the kinetic-energy square and exposes the transverse oscillator. A carrier-closed auxiliary field then preserves that transverse storage. A closed-to-open lift embeds a support point in the opened variables, after which full-field correction restores the original coupled locomotor dynamics. Only after a reduced motion has been found is the planar pose reconstructed. None of the first three operations is pose reconstruction, and carrier closure is neither a physical clamp nor the instruction to set one velocity to zero.

In the opened conservative locomotor, power may pass in either direction between the two reservoirs while their sum remains constant. The ideal constraint reaction performs no total work. Its admissible distribution and the constrained kinetic metric structure the coupling, but the constraint does not supply energy. The scalar theorem identifies the integrated exchange with the one return coordinate left unresolved on a positive oriented section.

Energy coordinates on regular scalar sections and periodic return maps are established tools ([Arnold, 1989](#); [Shaw and Pierre, 1993](#); [Kuznetsov, 2004](#)). The analytical contribution here is their locomotor connection: after the carrier coordinates return, zero net oscillator–channel exchange is necessary and sufficient for the remaining positive crossing speed to return. Twenty-nine finite-speed solutions realize this equivalence numerically and agree with the independently continued family of [Chapter 3](#). The result is exact under the stated section hypotheses, but it is not a general existence, stability, or multi-transverse return theorem.

4.1 The return coordinate to be replaced

Recall the pose-reduced state

$$z = (\delta, v, \omega, \sigma), \quad \sigma = \dot{\delta}, \quad (4.1)$$

where v is the body-forward speed and ω is the carrier yaw rate. The body pose $g = (x, y, \theta) \in \text{SE}(2)$ is reconstructed from

$$\dot{x} = v \cos \theta, \quad \dot{y} = v \sin \theta, \quad \dot{\theta} = \omega. \quad (4.2)$$




Figure 4.1: Energy and operation map of the closed–open construction. Carrier elimination exposes an effective transverse oscillator; the auxiliary closed field preserves its storage. Restoring the full field opens the propulsive channel, and power is exchanged with the displayed sign while total energy remains constant. The ideal constraint structures this coupling but supplies no power. Pose reconstruction is a later integration and is not shown as a closing or opening operation.

A relative-periodic locomotor motion closes z after a period T while allowing $g(T) = g(0)\Delta g$ with $\Delta g \neq e_G$. The group displacement is an output of the returned reduced motion; it is not included in the return residual.

A Poincaré formulation removes the arbitrary time origin of a periodic orbit (Kuznetsov, 2004). On the section $\delta = 0$, the first return has three nontrivial coordinates. The residual used here imposes direct return of the two carrier coordinates and replaces the crossing-speed defect by an exchange integral, after fixing the section orientation, storage, and carrier-coupled field.

Shooting and pseudo-arclength continuation find zeros of a supplied residual (Peeters et al., 2009). The analytical result below identifies the power-balance row with one coordinate of the section return for this scalar transverse problem.

4.2 Carrier elimination and the effective oscillator

The complete physical picture precedes the coefficient-level derivation in Figure 4.1. Carrier elimination identifies the two reservoirs by completing the kinetic-energy square. The auxiliary carrier-closed field retains the effective oscillator as its independent energy reservoir and conserves $E_\perp$. Restoring the full field reactivates the complementary propulsive storage. The resulting exchange changes each reservoir, not their sum.

4.2.1 Completed square and transverse storage

In admissible velocities $\eta = (v, \omega, \sigma)^\top$, the kinetic energy of the dimensionless two-segment model is

$$K(\delta, \eta) = \frac{1}{2} \eta^\top M(\delta) \eta, \quad (4.3)$$

with

$$M(\delta) = \begin{bmatrix} 1 & -\rho & -\rho \\ -\rho & B_{11} & B_{12} \\ -\rho & B_{12} & B_{22} \end{bmatrix}. \quad (4.4)$$

The functions $\rho(\delta)$, $B_{11}(\delta)$, and $B_{12}(\delta)$, and the constant B_{22} , were derived in [Chapter 3](#). For the reference design,

$$B_{11} = j_1 + j_2 + \alpha^2 + \mu\beta^2 + 2\mu\alpha\beta \cos \delta, \quad (4.5)$$

$$B_{12} = j_2 + \mu\beta^2 + \mu\alpha\beta \cos \delta, \quad (4.6)$$

$$B_{22} = j_2 + \mu\beta^2, \quad \rho = \mu\beta \sin \delta. \quad (4.7)$$

The complete dimensional derivation and parameter table are collected in [Chapter B](#).

Partition the velocities into the carrier pair $y = (v, \omega)^\top$ and the scalar shape rate σ :

$$M = \begin{bmatrix} M_{cc} & m_{c\sigma} \\ m_{c\sigma}^\top & B_{22} \end{bmatrix}, \quad M_{cc} = \begin{bmatrix} 1 & -\rho \\ -\rho & B_{11} \end{bmatrix}, \quad m_{c\sigma} = \begin{bmatrix} -\rho \\ B_{12} \end{bmatrix}. \quad (4.8)$$

When

$$\Delta(\delta) = \det M_{cc} = B_{11} - \rho^2 > 0, \quad (4.9)$$

the carrier block can be eliminated at fixed (δ, σ) . The energy-minimizing carrier velocity and its defect coordinate are

$$y_* = -M_{cc}^{-1} m_{c\sigma} \sigma, \quad q_c = y - y_* = y + M_{cc}^{-1} m_{c\sigma} \sigma. \quad (4.10)$$

The coordinate q_c measures departure from the carrier-closed leaf. It is a completed-square velocity coordinate; the argument does not require it to be a conserved pseudo-momentum. In particular, $q_c = 0$ means $y = y_*(\delta, \sigma)$, which is generally nonzero; closing this completed-square coordinate neither stops the carrier nor forces the reconstructed pose to close. The effective transverse inertia is

$$M_\perp(\delta) = B_{22} - m_{c\sigma}^\top M_{cc}^{-1} m_{c\sigma} = B_{22} - \frac{B_{12}^2 + \rho^2(B_{11} - 2B_{12})}{\Delta}. \quad (4.11)$$

Positive definiteness of the admissible kinetic energy implies $M_\perp > 0$ wherever M and M_{cc} are positive definite.

With the internal potential

$$U(\delta) = \frac{1}{2}k_2\delta^2 + \frac{1}{4}k_4\delta^4, \quad k_2 = 1, \quad k_4 = 10, \quad (4.12)$$

the scalar transverse storage is

$$E_\perp(\delta, \sigma) = \frac{1}{2}M_\perp(\delta)\sigma^2 + U(\delta). \quad (4.13)$$

Completing the square now gives the exact split

$$E = E_{\text{prop}} + E_\perp, \quad E_{\text{prop}} = \frac{1}{2}q_c^\top M_{cc}(\delta)q_c. \quad (4.14)$$

Thus $E_\perp$ is not the total energy of the locomotor. It is the kinetic-plus-potential storage of the effective transverse oscillator exposed by carrier elimination. The complementary completed-square term is the storage associated with the propulsive channel. Calling the channel “propulsive” identifies its carrier role in pose transport; it is not a causal decomposition of net displacement. More precisely, E_{prop} measures the energetic defect from the minimizing carrier velocity y_* , not all carrier kinetic energy: the Schur-complement term $E_\perp$ already includes the carrier motion induced by y_* . The split is therefore exact for the declared carrier–shape partition, not a unique coordinate-free attribution of propulsion.




Figure 4.2: The scalar geometry used by the return theorem. Regular closed levels foliate an annulus and supply an initial scale; the oriented section removes the half-period ambiguity and makes transverse storage one-to-one in the crossing speed.

4.2.2 The closed transverse field

The auxiliary carrier-closed field evolves (δ, σ) with $E_\perp$ conserved. Differentiating Equation (4.13) gives

$$\dot{E}_\perp = \sigma \left(M_\perp \dot{\sigma} + \frac{1}{2} M'_\perp \sigma^2 + U' \right). \quad (4.15)$$

The closed field is therefore

$$\dot{\delta} = \sigma, \quad \dot{\sigma} = F_{\text{cl}}(\delta, \sigma) = -\frac{U'(\delta) + \frac{1}{2} M'_\perp(\delta) \sigma^2}{M_\perp(\delta)}. \quad (4.16)$$

At a regular energy E , its two branches are

$$\sigma_\pm(\delta; E) = \pm \sqrt{\frac{2(E - U(\delta))}{M_\perp(\delta)}}, \quad (4.17)$$

between the turning points $U(\delta_\pm) = E$. Period and action follow from one-dimensional quadratures,

$$\frac{T_{\text{cl}}(E)}{2} = \int_{\delta_-}^{\delta_+} \sqrt{\frac{M_\perp(\delta)}{2(E - U(\delta))}} d\delta, \quad (4.18)$$

$$J_{\text{cl}}(E) = \frac{1}{\pi} \int_{\delta_-}^{\delta_+} \sqrt{2M_\perp(\delta)(E - U(\delta))} d\delta. \quad (4.19)$$

This is the precise pendulum–2SEG analogy. A simple pendulum has one configuration coordinate and its regular energy level fixes the two velocity branches. The complete 2SEG locomotor has four configuration coordinates and three admissible velocity coordinates, but carrier elimination exposes one effective transverse coordinate with the same scalar level-set geometry. The variable inertia $M_\perp(\delta)$ changes the clock and action; it does not change the one-dimensional character of the oriented crossing used below.

These quantities parameterize and seed the closed support. They do not yet describe an opened locomotor cycle. The regular level and the oriented crossing used below are shown in Figure 4.2.

4.3 Closed-to-open lift and restoration of the full field

The carrier-closed support lives in (δ, σ) . A *closed-to-open lift* augments a point on that support with carrier coordinates and therefore produces an initial candidate in the complete pose-reduced state $(\delta, v, \omega, \sigma)$. The candidate is then corrected under the original projected equations, so the propulsive channel is allowed to evolve rather than prescribed. This restoration of the full field is the operation called *reopening*. It remains distinct from the final pose reconstruction, which integrates Equation (3.6) only after the reduced return problem has been solved.

4.3.1 Restored finite-speed field

The full reduced dynamics couples shape acceleration to the carrier through the product

$$q_y = v\omega. \quad (4.20)$$

Solving the original reduced equations for $\dot{\sigma}$ gives the restored opened field,

$$F_{\text{op}}(\delta, \sigma, v, \omega) = \frac{r(\delta)[q_y - \rho(\delta)(\sigma + \omega)^2] - U'(\delta)}{M_{\perp}(\delta)}, \quad (4.21)$$

where

$$r(\delta) = \frac{A_q(\delta)}{\Delta(\delta)} \quad (4.22)$$

and A_q is the coefficient obtained when the carrier block is eliminated from the raw equations. Chapter C records the remaining transport rows and the equation-to-code regression.

The distinction between Equations (4.16) and (4.21) fixes the convention used throughout this thesis. In particular, the metric derivative is part of F_{cl} but is not inserted into F_{op} . The exchange force is the difference of the two fields:

$$M_{\perp}(F_{\text{op}} - F_{\text{cl}}) = r\{q_y - \rho(\sigma + \omega)^2\} + \frac{1}{2}M'_{\perp}\sigma^2. \quad (4.23)$$

Thus closing the channel means replacing the opened field by the conservative field. It is not equivalent to setting only q_y or ω to zero in Equation (4.21).

Multiplication by the shape rate identifies the instantaneous transverse exchange power,

$$P_{\text{ex}} = \sigma M_{\perp}(F_{\text{op}} - F_{\text{cl}}) = \sigma \left[r\{q_y - \rho(\sigma + \omega)^2\} + \frac{1}{2}M'_{\perp}\sigma^2 \right]. \quad (4.24)$$

Substitution of the opened field into the derivative of Equation (4.13) yields the exact identity

$$\dot{E}_{\perp} = P_{\text{ex}}, \quad \dot{E}_{\text{prop}} = -P_{\text{ex}}. \quad (4.25)$$

With this sign convention, $P_{\text{ex}} > 0$ means that the effective oscillator is receiving power from the propulsive channel; $P_{\text{ex}} < 0$ means that it is giving power to that channel. The second equality follows from the completed-square split Equation (4.14) and conservation of the total energy. The ideal reaction contributes no term to $\dot{E}$. The integrated exchange along any regular unforced conservative opened segment is consequently

$$I_{\text{ex}} = \int_0^T P_{\text{ex}}(t) \, dt = E_{\perp}(T) - E_{\perp}(0). \quad (4.26)$$

Because the transverse crossing-speed defect is omitted from the numerical residual, (4.26) supplies that scalar return row through the exchange balance.

4.3.2 Finite-speed and orientation chart

Let the prescribed nonzero mean speed be $\bar{v}$ and set

$$h = |\bar{v}|^{-1}, \quad s = \text{sgn}(\bar{v}), \quad \nu = hv, \quad qq_{yy} = v\omega. \quad (4.27)$$

The inverse relations are

$$v = \frac{\nu}{h}, \quad \omega = \frac{hq_y}{\nu}, \quad (4.28)$$

so the chart requires $\nu \neq 0$. It keeps q_y finite as $h \rightarrow 0$ and separates the large forward speed from the smaller yaw rate.

Positive and negative mean-speed branches are placed on one oriented sheet with

$$Y = (\delta, \zeta, u, Q), \quad \sigma = s\zeta, \quad \nu = su, \quad q_y = sQ, \quad (4.29)$$

and oriented time ϑ defined by $d/d\vartheta = s d/dt$. Then $d\delta/d\vartheta = \zeta$, and both speed directions use the section

$$\Sigma = \{Y : \delta = 0, \zeta > 0, u > 0\}. \quad (4.30)$$

The full oriented field has the form

$$\frac{dY}{d\vartheta} = (\zeta, F_{\text{op}}, F_u, F_Q), \quad (4.31)$$

where F_u and F_Q are algebraic transforms of the raw carrier equations. Their equation-to-code comparison is reported in [Chapter C](#).

4.4 Exact scalar exchange–return equivalence

Let $Y_0 = (0, \zeta_0, u_0, Q_0) \in \Sigma$ and let

$$Y^+ = (0, \zeta^+, u^+, Q^+) \quad (4.32)$$

be the next crossing of the same oriented section. The section storage depends on one coordinate:

$$E_\Sigma(\zeta) = E_\perp(0, s\zeta) = \frac{1}{2}M_\perp(0)\zeta^2 + U(0). \quad (4.33)$$

Since $M_\perp(0) > 0$ and $\zeta > 0$,

$$\frac{dE_\Sigma}{d\zeta} = M_\perp(0)\zeta > 0. \quad (4.34)$$

Orientation is essential: without it, the two crossing directions $\pm\zeta$ would have the same storage.

Physical statement. Suppose a trajectory leaves $\delta = 0$ with positive oriented crossing coordinate ζ and returns to the same oriented section after its two carrier coordinates have returned. There, transverse storage increases strictly with ζ . Hence the effective oscillator has zero net energy gain from the propulsive channel if and only if its oriented crossing coordinate returns. The physical joint rate is $\sigma = s\zeta$, so its sign is reversed on the negative-mean-speed sheet. Zero integrated exchange therefore supplies the missing scalar return condition; it is not an extra test on an orbit whose complete state has already returned.

Assumption 4.1 (minimal regular scalar return domain). The trajectory follows the unforced conservative opened field and the storage identity [Equation \(4.25\)](#). It remains in a domain on which $\Delta > 0$, $M_\perp > 0$, u is bounded away from zero, and the finite-speed field is regular. The flow crosses [Equation \(4.30\)](#) transversely and has a well-defined next return on the same oriented sheet.

Theorem 4.2 (scalar exchange–return equivalence). *Under ?? 4.1, the integrated exchange and the transverse crossing-speed defect are equivalent:*

$$I_{\text{ex}} = 0 \iff \zeta^+ = \zeta_0. \quad (4.35)$$

Consequently, if $u^+ = u_0$ and $Q^+ = Q_0$ are imposed directly, the residual row $I_{\text{ex}} = 0$ completes the pose-reduced return on Σ .

Proof. By Equation (4.26), $I_{\text{ex}} = 0$ is equivalent to $E_{\perp}(T) = E_{\perp}(0)$. Both endpoints lie at $\delta = 0$ on the same oriented sheet, hence

$$E_{\Sigma}(\zeta^+) = E_{\Sigma}(\zeta_0). \quad (4.36)$$

The strict monotonicity of Equation (4.33) on $\zeta > 0$ gives $\zeta^+ = \zeta_0$. Conversely, equality of the crossing speeds makes the endpoint transverse storages equal, so Equation (4.26) gives $I_{\text{ex}} = 0$. $\square$

The theorem identifies the zero set of one return row; it does not by itself prove existence or stability. For the 29 roots below, the scaled return-Jacobian analysis establishes local regularity at the sampled points. The three return rows have numerical rank two; their one differential dependence is required by the total-energy first integral. Adding the mean-speed row gives a full-column-rank 4×3 Jacobian at every fixed- $\bar{v}$ root, with smallest singular value in $[0.021917, 0.104764]$. When $\bar{v}$ is restored as a branch coordinate, the 4×4 Jacobian has rank three: its smallest nonzero singular value lies in $[0.025463, 0.107124]$, while the largest null singular value is 1.64×10^{-10} . Thus every tested root is a regular sampled point of a local phase-fixed branch. For a root with nonidentity reconstructed pose, restoring cycle phase gives a local NLM chart wherever the sweep is embedded. This certification is local to the sampled nonzero-speed sheets and does not join the roots through $\bar{v} = 0$ or establish global regularity.

On a higher-dimensional section, equality of one storage value defines a level set rather than identifying the endpoint. Full vector return must then be imposed directly, and modal metrics separately classify the returned cycle.

4.5 A computable exchange residual

For fixed $\bar{v} \neq 0$, write the initial section point as

$$Y_0 = (0, e^a, e^b, Q_0). \quad (4.37)$$

The logarithmic coordinates enforce $\zeta_0 > 0$ and $u_0 > 0$. The opened field is integrated to the next event Equation (4.30), together with two accumulators,

$$\frac{dI_{\text{ex}}}{d\vartheta} = sP_{\text{ex}}, \quad \frac{dI_u}{d\vartheta} = u. \quad (4.38)$$

The factor s expresses physical-time exchange power in oriented time. For $s = -1$ the numerical orbit represents a backward physical-time traversal; its cycle integral changes sign relative to the forward traversal but has the same zero set. The overdetermined shooting residual is

$$\mathcal{R}_{\text{ex}}(Y_0; \bar{v}) = \begin{bmatrix} I_{\text{ex}}(T) \\ u(T) - u_0 \\ Q(T) - Q_0 \\ T^{-1}I_u(T) - 1 \end{bmatrix}. \quad (4.39)$$

The first three rows close the reduced section state by Theorem 4.2; the fourth fixes the mean-speed gauge introduced by $h = |\bar{v}|^{-1}$. Although four components are reported for three initial coordinates,

Table 4.1: Representative exchange-return rows and comparison with the independent continued family. Errors are absolute; the geometry column is a scaled root-mean-square (RMS) distance after cyclic phase and traversal alignment.

$\bar{v}$	$\ \mathcal{R}_{\text{ex}}\ $	period error	δ_{max} error	geometry RMS
-24.871	1.80×10^{-15}	4.72×10^{-6}	9.19×10^{-7}	3.07×10^{-3}
-10.162	9.72×10^{-15}	6.64×10^{-7}	5.13×10^{-6}	2.04×10^{-3}
-0.897	5.52×10^{-13}	3.85×10^{-7}	1.99×10^{-6}	1.90×10^{-3}
1.133	2.99×10^{-15}	6.02×10^{-7}	3.55×10^{-6}	1.99×10^{-3}
10.274	6.38×10^{-14}	8.19×10^{-6}	6.64×10^{-6}	2.06×10^{-3}
25.088	1.76×10^{-15}	1.50×10^{-5}	1.86×10^{-5}	2.81×10^{-3}

they are not four independent conditions. Energy preservation forces one differential dependence among the three section-return rows; the Jacobian analysis resolves the remaining rank as two and shows that the mean-speed row supplies the third independent direction. The resulting 4×3 Jacobian has full column rank at all 29 tested roots. All four scaled rows are minimized simultaneously by nonlinear least squares and pass their common tolerance; [Chapter C](#) records the scaling and rank calculation.

In the computation, a regular closed-support annulus without separatrix or turning-point singularity supplies the initial scale. This additional support condition is numerical and is not required by the equivalence proof itself. The propulsive channel is restored, the opened field is integrated to the next oriented crossing, and the exchange and two carrier-return rows are solved simultaneously. Acceptance requires the section event, the regular-domain checks, and all four scaled residual rows to pass. Comparison with the continuation of [Chapter 3](#) is then an independent test of branch placement, not part of the residual definition.

4.6 Numerical realization

4.6.1 Returned family

The numerical data contain 29 accepted rows spanning both signs of the central finite-speed corridor. Every row closes [Equation \(4.39\)](#); the largest scaled residual norm is

$$5.51 \times 10^{-12}. \quad (4.40)$$

The largest mismatch between the requested $\bar{v}$ and the nearest independently continued reference speed is 7.35×10^{-11} . Representative rows are listed in [Table 4.1](#); aggregate comparison bounds and independent event/domain checks are provided in [Chapter C](#).

Across all 29 rows, the maximum absolute errors are 1.50×10^{-5} for period, 1.86×10^{-5} for maximum shape amplitude, and 8.36×10^{-4} for maximum yaw-rate amplitude. The maximum phase-aligned scaled RMS loop distance is 3.07×10^{-3} . Time-sample differences are not used because two representations of the same periodic orbit may choose different origins and traversal directions. [Figures 4.3](#) and [4.4](#) show, respectively, the residual decomposition and representative phase-aligned loop comparisons.

4.7 Dimensional boundary of the scalar result

The closed–open operation itself extends beyond one transverse coordinate: one can eliminate carrier directions, construct a reduced storage or modal support, restore the carrier coupling, and evaluate




Figure 4.3: Closure of the exchange and carrier-return rows over the 29 accepted cycles. The left panel reports the scaled residual norm; the right panel separates its four absolute components.




Figure 4.4: Exchange–return cycles (solid colored curves) and independently continued reference loops (black dashed curves) at representative mean speeds. Agreement is evaluated after cyclic-phase and traversal alignment.

return and exchange. The equivalence in [Theorem 4.2](#) does not extend by notation alone. Its decisive step is the injective map

$$\zeta \mapsto E_{\Sigma}(\zeta) \quad (4.41)$$

on a one-dimensional oriented section.

For a transverse state $w \in \mathbb{R}^m$ with $m > 1$, the equation $E_{\perp}(w^+) = E_{\perp}(w_0)$ generally leaves an $(m - 1)$ -dimensional set of possible endpoints. Exchange still measures storage transfer between transverse and carrier directions. Full vector return then establishes exact reduced return, while modal metrics distinguish the returned organization in the three-segment locomotor.

Result and dimensional boundary

Carrier elimination exposes the effective transverse oscillator and the complementary propulsive storage by an exact completed square. The carrier-closed field is an auxiliary conservative dynamics for that oscillator; the closed-to-open lift and full-field correction restore the original coupled locomotor equations. Along the restored field, $P_{\text{ex}} = \dot{E}_{\perp} = -\dot{E}_{\text{prop}}$: the two reservoirs exchange power while the ideal constraint supplies none.

On the regular positive scalar section, $E_{\perp}$ is injective in the remaining crossing speed. Zero integrated exchange can therefore replace that final return row after the carrier coordinates return. The equivalence is exact, but its decisive step is one-dimensional. With several transverse coordinates, equality of one storage value leaves a level set of possible endpoints. [Chapter 5](#) therefore keeps vector return and uses exchange only after that return has been established.

Chapter 5

Modal families of the three-segment locomotor

The *three-segment locomotor* (3SEG) adds a second internal joint to the same constrained carrier architecture. On the declared regular one-event section, absent additional independent integrals or symmetries, this apparently small change crosses the exact boundary of the scalar theorem. The effective transverse oscillator now has two coordinates, just as the double pendulum has more than one way to organize a given total energy. One scalar storage balance can restrict the endpoint, but it cannot identify the endpoint or its modal organization.

The calculation therefore keeps three questions in a fixed order. First, all six pose-reduced state components must return under the unforced field. Second, the already-returned motion is classified by its finite-amplitude modal content. Third, energy exchange and pose are evaluated as consequences: exchange checks the transverse-storage balance, while pose reconstruction determines the rigid-body displacement. Neither modal identity nor nonzero transport substitutes for vector return.

At one fixed dimensionless model-parameter vector p_* , this procedure yields two moving branch chains continued from the two linear modal directions. The parent modal seeds are termed *in-phase* (IP) and *anti-phase* (AP). Their finite-amplitude members are described as *in-phase-descended* and *anti-phase-descended*: continuation preserves a recognizable modal organization, but does not imply exact zero or π phase locking or a purely sinusoidal waveform. A matched straight SE(2) motion then separates the pose increment of each cycle from a reference with the same mean speed and period. The result is numerical coexistence of two modal families of the same locomotor, not stability, exhaustive branch classification, or one global modal manifold.

5.1 Why one storage value cannot identify the endpoint

After the carrier rows have returned, let w collect the unresolved coordinates on an oriented section. In the scalar case, $w = \zeta > 0$, and strict monotonicity of the section storage makes endpoint storage injective. For $w \in \mathbb{R}^m$, $m > 1$,

$$E_\Sigma(w^+) = E_\Sigma(w_0) \implies w^+ \in E_\Sigma^{-1}(E_\Sigma(w_0)), \quad (5.1)$$

and a regular level set is generally $(m - 1)$ -dimensional. Equality of one scalar storage therefore constrains the endpoint but does not determine it.

For the actual 3SEG state $z \in \mathbb{R}^6$, consider for instance the oriented section $\delta_1 = 0$. Removing that event coordinate and returning the two carrier velocities still leaves

$$w = (\delta_2, \sigma_1, \sigma_2) \in \mathbb{R}^3. \quad (5.2)$$

A regular scalar storage equality therefore leaves a two-dimensional surface of possible section

Table 5.1: Exact boundary for the declared regular one-event problem between the scalar 2SEG result and the modal 3SEG construction, absent additional independent integrals or symmetries. In both cases the propulsive channel and effective oscillator exchange power; only the dimension of the missing return information changes.

Question	2SEG locomotor	3SEG locomotor
effective transverse degrees of freedom	one: coordinate δ , state (δ, σ)	two: coordinates (δ_1, δ_2) and their rates
oriented storage level	two crossing directions; orientation leaves one point per value	a positive-dimensional level set remains; it is a surface for the section in Equation (5.2)
exact qualification	carrier return plus scalar exchange–return theorem	six-component reduced vector return
role of exchange	replaces the final scalar return row	endpoint identity and numerical balance check after return
modal information	no competing transverse direction on the scalar section	in-phase- or anti-phase-descended identity evaluated after return
local family geometry	one phase-fixed branch under full-rank regularity	separate sector branches under full-rank regularity

endpoints. The same codimension count holds for any regular one-event section; its coordinates may change, but one energy value cannot close the remaining vector state.

The following ellipse is a two-dimensional slice illustrating that non-uniqueness:

$$E_{\Sigma}(w_1, w_2) = \frac{1}{2} (\kappa_1 w_1^2 + \kappa_2 w_2^2), \quad \kappa_1, \kappa_2 > 0, \quad (5.3)$$

assigns the same value to $(a, 0)$ and $(0, a\sqrt{\kappa_1/\kappa_2})$, although the points lie on different modal axes. Choosing one crossing orientation removes a sign ambiguity; it does not collapse the ellipse to a point.

This dimensional obstruction fixes the order of the chapter’s claims. Full vector return first qualifies a reduced cycle. Modal metrics subsequently classify the returned motion as in-phase- or anti-phase-descended. Once return is known, equality of endpoint transverse storage makes zero cycle-mean exchange an identity and a numerical consistency check, not a third independent selector.

Table 5.1 places the scalar theorem and the present vector problem face to face. The distinction is not that exchange ceases to matter; it is that one exchanged-energy balance ceases to contain enough coordinates to certify return.

5.2 Locomotor and reduced conservative field

5.2.1 Configuration, state, and pose reconstruction

The 3SEG locomotor consists of a constrained carrier, a middle segment, and a distal segment. Its two relative angles are $\delta = (\delta_1, \delta_2)^T$: δ_1 is measured from the carrier to the middle segment and δ_2 from the middle to the distal segment. Figure 5.1 fixes the associated velocity and angle conventions before the reduced equations are introduced.

The configuration, admissible velocity, and pose-reduced state are

$$q = (g, \delta), \quad g = (x, y, \theta) \in \text{SE}(2), \quad \xi = (v, \omega, \sigma_1, \sigma_2)^T, \quad (5.4)$$

$$z = (\delta_1, \delta_2, v, \omega, \sigma_1, \sigma_2)^T. \quad (5.5)$$

The no-side-slip constraint and an admissible-velocity basis are

$$A_3(\theta)\dot{q} = 0, \quad A_3(\theta) = \begin{bmatrix} -\sin \theta & \cos \theta & 0 & 0 & 0 \end{bmatrix}, \quad \dot{q} = N(q)\xi, \quad (5.6)$$




Figure 5.1: Coordinates of the 3SEG locomotor. Each joint angle is relative to the preceding segment. The carrier moves forward at speed v , yaws at rate ω , and cannot slip laterally. Complete centre-of-mass definitions are given in Chapter D.

with

$$N(q) = \begin{bmatrix} \cos \theta & 0 & 0 & 0 \\ \sin \theta & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \text{Im } N(q) = \ker A_3(\theta). \quad (5.7)$$

This is exactly the pattern of Equation (3.4) and Table 3.1, with N in the role of H . The configuration has five coordinates, the rank-one constraint leaves four admissible velocities, two are carrier velocities (v, ω) , and two are effective transverse joint velocities (σ_1, σ_2) . Projection gives

$$N(q)^\top \left[\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q \right] = 0. \quad (5.8)$$

The same smooth, ideal, constant-rank scope applies; the extra segment changes the dimension and modal organization, not the reduction method.

Once $z(t)$ has been computed, the planar pose is reconstructed from

$$\dot{x} = v \cos \theta, \quad \dot{y} = v \sin \theta, \quad \dot{\theta} = \omega. \quad (5.9)$$

Return is imposed on all six components of z ; the pose is left open and reported through $\Delta g = g(0)^{-1}g(T)$.

5.2.2 Unforced equations and total energy

The joints have local quadratic–quartic springs and a weak quadratic coupling,

$$U(\delta) = \frac{1}{2}k_{2,1}\delta_1^2 + \frac{1}{4}k_{4,1}\delta_1^4 + \frac{1}{2}k_{2,2}\delta_2^2 + \frac{1}{4}k_{4,2}\delta_2^4 + \frac{1}{2}k_{12}(\delta_1 - \delta_2)^2. \quad (5.10)$$

With $\sigma = (\sigma_1, \sigma_2)^\top$ and $\eta = (v, \omega)^\top$, the unforced reduced equations are

$$\dot{\delta} = \sigma, \quad M(\delta) \begin{bmatrix} \dot{\eta} \\ \dot{\sigma} \end{bmatrix} + c(\delta, \eta, \sigma) + \begin{bmatrix} 0 \\ 0 \\ \nabla U(\delta) \end{bmatrix} = 0. \quad (5.11)$$

The exact geometry and kinetic-energy definition of the mass matrix, together with the velocity-dependent vector, are given in Chapter D. The ideal constraint does no total work, so

$$E(z) = \frac{1}{2}\xi^\top M(\delta)\xi + U(\delta) \quad (5.12)$$

is conserved by this unforced field. The numerical parameter values used below define one fixed model and carry no assigned physical units.

5.3 Closed supports, modal seeds, and the closed-to-open lift

5.3.1 Carrier-closed transverse storage

Partition the mass matrix between carrier velocity η and joint velocity σ :

$$M = \begin{bmatrix} M_{GG} & M_{GS} \\ M_{SG} & M_{SS} \end{bmatrix}. \quad (5.13)$$

The subscripts G and S denote the pose-group carrier block and the shape block, respectively; they do not introduce additional coordinates. The energy-minimizing carrier velocity and complementary carrier defect are

$$\eta_{\perp} = -M_{GG}^{-1}M_{GS}\sigma, \quad q_c = \eta - \eta_{\perp}. \quad (5.14)$$

Eliminating that carrier velocity gives the Schur-complement inertia and transverse storage

$$M_{\perp} = M_{SS} - M_{SG}M_{GG}^{-1}M_{GS}, \quad E_{\perp}(\delta, \sigma) = \frac{1}{2}\sigma^{\top}M_{\perp}(\delta)\sigma + U(\delta). \quad (5.15)$$

On the domain used here, $M_{\perp}$ is positive definite. As in the scalar locomotor, completing the kinetic-energy square gives

$$E = E_{\text{prop}} + E_{\perp}, \quad E_{\text{prop}} = \frac{1}{2}q_c^{\top}M_{GG}(\delta)q_c. \quad (5.16)$$

Here E_{prop} measures the energetic defect from the minimizing carrier velocity $\eta_{\perp}$, not all carrier kinetic energy. Thus $q_c = 0$ still gives $\eta = \eta_{\perp}(\delta, \sigma)$, which is generally nonzero, and does not by itself imply zero pose transport.

The auxiliary field actually used to compute the closed supports is the Euler–Lagrange field of

$$L_{\perp}(\delta, \sigma) = \frac{1}{2}\sigma^{\top}M_{\perp}(\delta)\sigma - U(\delta). \quad (5.17)$$

Writing $D_i M_{\perp} = \partial M_{\perp} / \partial \delta_i$, its implemented form is

$$\begin{aligned} \dot{\delta} &= \sigma, \\ M_{\perp} \dot{\sigma} &= \frac{1}{2} \begin{bmatrix} \sigma^{\top} (D_1 M_{\perp}) \sigma \\ \sigma^{\top} (D_2 M_{\perp}) \sigma \end{bmatrix} - (\sigma_1 D_1 M_{\perp} + \sigma_2 D_2 M_{\perp}) \sigma - \nabla U. \end{aligned} \quad (5.18)$$

It satisfies $\dot{E}_{\perp} = 0$ exactly. Energy conservation alone would not define a two-coordinate field; [Equations \(5.17\) and \(5.18\)](#) specify the particular auxiliary dynamics used here. Its periodic solutions are called closed supports.

The relation $q_c = 0$ is then used to lift a support into the opened variables; it is not asserted to be an invariant restriction of the full nonholonomic field. Full-field correction lets q_c evolve. Along the resulting unforced conservative motion, $\dot{E}_{\perp} = -\dot{E}_{\text{prop}}$ while total energy remains constant. The ideal constraint supplies no power; its admissible distribution and the kinetic metric structure the exchange.

5.3.2 Local in-phase and anti-phase seeds

At the straight configuration, the mass-normalized eigenproblem is

$$K_0 u_j = \Omega_j^2 M_0 u_j, \quad u_i^{\top} M_0 u_j = \delta_{ij}, \quad j \in \{\text{IP}, \text{AP}\}, \quad (5.19)$$

with $M_0 = M_{\perp}(0)$ and $K_0 = D^2 U(0)$. The entries of u_{IP} have the same sign; those of u_{AP} have opposite signs. These eigenvectors provide local tangent directions for the closed-support families, not finite-amplitude trajectories. Continuation is needed because period, waveform, and modal mixture change away from the straight configuration ([Kerschen et al., 2009](#); [Peeters et al., 2009](#)).

5.3.3 Closed-to-open lift and full-field correction

A closed support is embedded in the opened state using the energy-minimizing carrier velocity in Equation (5.14), initially with $q_c = 0$. This *closed-to-open lift* supplies a candidate, not a solution of the opened problem. The subsequent correction allows q_c to vary and solves the full field Equation (5.11). Here $q_c = 0$ is a coordinate leaf used by the lift, not an invariant leaf claimed for the opened field; it is neither a physical clamp nor the instruction to set a world velocity to zero. The word *reconstruction* remains reserved for integrating the planar pose after the opened reduced cycle has returned.

The numerical construction followed

$$\begin{aligned} \text{linear seed} &\longrightarrow \text{closed support} \longrightarrow \text{bridged candidate} \\ &\longrightarrow \text{corrected opened cycle} \longrightarrow \text{fixed-parameter branch.} \end{aligned} \quad (5.20)$$

Multiple shooting and collocation corrected the opened candidates, and pseudo-arclength continuation organized neighboring cycles. This is the construction route represented by the available results. The final cycle data support full-field, return, modal, exchange, and pose checks, but do not retain the closed-support projection for every member. A one-to-one genealogy from one particular support through the lift to every final cycle is therefore not claimed.

5.4 Opened cycles: vector return, modal classification, and exchange

5.4.1 Qualification by the full reduced field and vector return

Let $\Phi_{p_*}^T$ be the flow of the opened unforced field at the common dimensionless model-parameter vector p_* . A corrected relative-periodic cycle satisfies the six-component residual

$$R_z(z_0, T) = \Phi_{p_*}^T(z_0) - z_0 = 0 \in \mathbb{R}^6. \quad (5.21)$$

Pose does not appear in this equality; it is reconstructed afterwards. Full vector return is the exact reduced-return statement. Neither modal identity nor exchange replaces a component of R_z in this model.

For local family geometry, the independent vector-return equations and a phase condition can be collected as $F_j(u, \lambda) = 0$ within one declared modal sector j . If the phase-fixed Jacobian has its expected full rank, its local zero set is a one-dimensional branch; on the transported subset, restoring cycle phase gives a local NLM chart of that fixed six-dimensional reduced field. This statement is conditional because the stored population does not include a branch-wide rank analysis. It also applies separately to the in-phase- and anti-phase-descended components: their union is not asserted to be one smooth manifold through unexamined branch points or modal mixing.

Dynamic consistency is checked independently on every stored time interval, including the interval that returns to the initial state. Starting from the stored node z_k , two fourth-order Runge-Kutta half steps of the exact field predict $\hat{z}_{k+1}$. The population-wide propagation check reports

$$\epsilon_{\text{flow}} = \max_{\text{cycles}, k} \|\hat{z}_{k+1} - z_{k+1}\|_{\infty}, \quad \epsilon_{\text{flow}, \text{sc}} = \max_{\text{cycles}, k, i} \frac{|\hat{z}_{k+1, i} - z_{k+1, i}|}{\max\{1, \text{ptp}(z_i)\}}. \quad (5.22)$$

A one-step/two-half-step Richardson estimate checks integration error, and the worst interval in each sector is recomputed with tight-tolerance DOP853. This seam-inclusive local multiple-shooting test distinguishes a duplicated closed polyline from a trace that follows the declared field. It also avoids using a single long integration as the sole test on possibly unstable periodic orbits, where small export errors can be strongly amplified. Its reported maxima are mesh-dependent interval defects, not a mesh-independent norm of field error; the stored interval count and the two refinement checks provide their scale.

5.4.2 Classification after return

The linear eigenvectors define modal coordinates on each complete finite-amplitude cycle,

$$Q_j(t) = u_j^\top M_0 \delta(t), \quad \tilde{P}_j(t) = \frac{u_j^\top M_0 \sigma(t)}{\Omega_j}. \quad (5.23)$$

Their time-weighted activity is

$$A_j = \left[\frac{1}{T} \int_0^T \left(Q_j(t)^2 + \tilde{P}_j(t)^2 \right) dt \right]^{1/2}. \quad (5.24)$$

The classification uses the expected-sector share $A_j/(A_{IP} + A_{AP})$, the expected/opposite activity ratio, and time-weighted correlations of the two joint positions and rates. Positive correlations identify the in-phase sense; negative correlations identify the anti-phase sense. The thresholds and exact modal vectors are listed in [Chapter D](#).

These metrics describe the organization of a cycle that has already returned. They do not supply exact return, and they do not prove the existence of a global invariant modal manifold.

5.4.3 Exchange after full return

Along any trajectory of the opened unforced conservative field,

$$P_{\text{ex}} = \frac{dE_\perp}{dt} = -\frac{dE_{\text{prop}}}{dt}, \quad \bar{P}_{\text{ex}} = \frac{1}{T} \int_0^T P_{\text{ex}} dt = \frac{E_\perp(T) - E_\perp(0)}{T}. \quad (5.25)$$

Positive and negative instantaneous values record redistribution between the carrier and transverse storage under the derived decomposition. If the full state returns, then $(\delta(T), \sigma(T)) = (\delta(0), \sigma(0))$, so $\bar{P}_{\text{ex}} = 0$ follows from the endpoint identity. Its reported numerical value tests storage consistency; it is not an independent qualification or classification row.

The cycle-mean body-forward speed,

$$\bar{v} = \frac{1}{T} \int_0^T v(t) dt, \quad (5.26)$$

is used below to compare the two modal sectors and to construct a matched straight pose reference.

5.4.4 Following one cycle from modal support to locomotion

The role of each operation is clearest when one candidate is followed from its modal support to its physical outputs. As shown in [Figure 5.2](#), the closed support and lift generate an initial guess. Full-field correction and six-component return qualify the opened reduced cycle. Modal content is then classified, exchange is checked as an energy identity, and pose is reconstructed as an output.

5.5 Coexisting anti-phase- and in-phase-descended families

At the same numerical model-parameter vector $p_\star$, correction under the full unforced field yields both modal organizations. The represented branches contain

$$N_{AP} = 2278, \quad N_{IP} = 1358 \quad (5.27)$$

opened cycles. The anti-phase-descended chain has 1,139 rows on each signed continuation side; the in-phase-descended chain has 679 on each side. These counts measure sampling density along two




Figure 5.2: From transverse modal support to locomotor output. Dashed boxes generate a candidate; solid boxes qualify and interpret the corrected opened motion. Full vector return precedes both modal classification and pose reconstruction. Exchange is a consequence and balance check after return, not a substitute for it.

branches, not independent replications or proof of completeness. The stored period is accepted as a return time; a separate minimal-period or primitivity claim is not needed for the coexistence result and is not made here.

Endpoint return, time-weighted modal identity, and pose reconstruction were recomputed on every represented member. In addition, every one of the 3,433,152 stored intervals, closure intervals included, was propagated locally with the independently implemented exact field. Energy, exchange, and the earlier spline differentiation diagnostic supply complementary checks. The resulting quantities are summarized in Table 5.2; the verification procedure and parameter vector are given in Chapter D. The available evidence concerns the corrected opened cycles. As stated above, it does not recover the particular closed-support genealogy of every member.

Table 5.2: Computed properties of the two opened modal families at the same numerical parameter vector p_* . The local-flow rows evaluate every stored interval against the exact field, including cycle closure. Endpoint zeros mean equality at the stored numerical precision.

Quantity	anti-phase-descended	in-phase-descended
represented returned cycles	2,278	1,358
maximum endpoint defect	0	0
maximum local-flow defect, raw	1.432×10^{-7}	5.557×10^{-6}
maximum local-flow defect, scaled	1.001×10^{-8}	1.074×10^{-6}
maximum relative total-energy span	7.131×10^{-9}	5.948×10^{-9}
members meeting the time-weighted sector criterion	2,278/2,278	1,358/1,358
maximum periodic-spline field defect	1.102×10^{-3}	1.502×10^{-3}
maximum $ \bar{P}_{\text{ex}} $	1.692×10^{-8}	2.764×10^{-8}
reconstructed displacement-magnitude range	$[3.038 \times 10^{-3}, 79.923]$	$[5.563 \times 10^{-2}, 143.562]$
maximum reconstructed $ \Delta\theta $	5.96×10^{-9}	2.48×10^{-7}
mean-speed span	$[-126.956, 122.793]$	$[-155.563, 166.114]$

The result establishes coexistence of returned, moving anti-phase- and in-phase-descended branch chains at p_* . All represented cycles meet their declared sector thresholds, their seam-inclusive local-flow defects remain at most 1.074×10^{-6} after coordinate scaling, and their reconstructed pose increments are nonzero at the resolution reported. This establishes finite numerical coexistence; stability and exhaustive branch classification remain separate questions.

5.6 Finite-amplitude contrast between the sectors

A fixed display sample of 72 cycles from each sector covers both speed signs and several scales of $|\bar{v}|$. It is an illustration of the computed branches, not a population-weighted statistic. The coordinate-excursion indicator is the largest peak-to-peak range among the dimensionless stored coordinates $\delta_1, \delta_2, \omega, \sigma_1, \sigma_2$, excluding forward speed. Because it combines angles and rates, it is a descriptive display statistic, not a norm or an energy. Table 5.3 then compares period, coordinate excursion, modal share, and joint coordination on this common sample.

Table 5.3: Finite-amplitude contrast in the fixed 72-cycle display sample from each sector.

Quantity	anti-phase sample	in-phase sample
period	0.5563–0.6295	0.8452–0.8642
largest internal-coordinate peak-to-peak range	52.23–84.43	4.915–5.388
time-weighted expected-sector share	0.6925–0.8136	0.9633–0.9670
expected/opposite activity ratio	2.252–4.364	26.284–29.275
joint-position correlation	–0.99991––0.99959	0.9147–0.9293
joint-rate correlation	–0.99915––0.99630	0.5642–0.6141

The anti-phase-descended family is the shorter-period, larger-excursion organization. Its two joints remain almost perfectly opposed, while the finite-amplitude motion has more activity in the opposite linear coordinate than the in-phase-descended family does. The in-phase-descended family is longer-period, lower-activity, and more concentrated near its same-direction linear seed. These statements concern modal ancestry and full-cycle activity; they do not assert exact 0 or π phase locking at finite amplitude. Representative traces make the distinction visible. Figure 5.3 compares one returned cycle from each sector in normalized joint coordinates.

Shape traces alone do not yet show why these returned motions are locomotor cycles. Figure 5.4 reconstructs the planar locomotor for one member of each family near the same mean speed. Reading it together with Figure 5.3 connects internal choreography to body motion; the matched straight reference is introduced immediately afterwards.

5.7 Matched-straight pose decomposition

A nonzero pose increment shows that a returned cycle moves the locomotor, but it does not isolate the effect of its internal oscillation. A carrier with the same nonzero mean speed would already translate during the same period. The matched-straight comparison separates these two group increments without modifying the computed reduced trajectory.

For every cycle, pose is reconstructed from Equation (5.9) with $g(0) = e_G$. The counterfactual is the straight relative equilibrium with the same sampled time-weighted mean speed and period:

$$\xi_{\text{base}} = [\bar{v} \ 0 \ 0]^T, \quad \Delta g_{\text{base}} = \text{Exp}_{\text{SE}(2)}\left(T\hat{\xi}_{\text{base}}\right), \quad \Delta g_{\text{inc}} = \Delta g_{\text{base}}^{-1} \Delta g_{\text{cycle}}. \quad (5.28)$$

The multiplication order expresses the residual in the carrier frame at the beginning of the cycle. Let

$$\text{Log}_{\text{SE}(2)}(\Delta g_{\text{inc}})^\vee = (\rho_{\text{inc}}, \phi_{\text{inc}}), \quad \rho_{\text{inc}} \in \mathbb{R}^2, \quad (5.29)$$

where the principal logarithm is guarded against rotational windings. Every represented 3SEG cycle lies on its reliable principal branch.

The normalized quantity

$$r_{\text{inc}} = 100 \frac{\|\rho_{\text{inc}}\|}{|\bar{v}| T} \quad (\%) \quad (5.30)$$




Figure 5.3: Finite-amplitude choreography of two returned 3SEG cycles near $\bar{v} = -8$. The top row compares the two joint angles over one period, normalized by their respective peak-to-peak excursions; the annotations give the dimensionless period and the unnormalized excursions. The bottom row shows the corresponding joint-space loops. The anti-phase-descended member retains opposed coordination and the in-phase-descended member retains same-direction coordination, although neither waveform is required to keep an exact linear phase relation. These examples identify modal ancestry, not stability or attraction.




Figure 5.4: Common-duration pose sequences from representative cycles of the same 3SEG locomotor near $\bar{v} = -8$. The anti-phase-descended cycle has $\bar{v} = -7.994$ and $T = 0.562$; the in-phase-descended cycle has $\bar{v} = -8.095$ and $T = 0.853$. The common displayed interval is one in-phase-descended period and approximately 1.52 anti-phase-descended periods, so the anti-phase trace passes its first return before the displayed endpoint. The displayed longitudinal increments are 6.906 for the in-phase-descended motion and 6.820 for the anti-phase-descended motion. They are increments over this common interval, not a comparison of two one-cycle increments. Intermediate outlines reveal the different joint choreographies. These are dimensionless model reconstructions from returned cycles, not hardware trajectories or evidence of stability.

is reported only for $|\bar{v}| \geq 1$. This threshold is a numerical reporting guard: as the matched straight displacement tends to zero, the ratio becomes ill-conditioned. The absolute residual $\|\rho_{\text{inc}}\|$ remains defined for every cycle, including the 58 anti-phase-descended and 8 in-phase-descended rows inside the guard.

Figure 5.5 shows the resulting absolute and guarded relative residuals for both complete populations.




Figure 5.5: Matched-straight pose decomposition for all returned anti-phase- and in-phase-descended cycles. Panel (a) compares reconstructed longitudinal displacement with the straight reference $\bar{v}T$. Panel (b) reports the translational component $\|\rho_{\text{inc}}\|$ of the principal SE(2) logarithm of Δg_{inc} . Panel (c) normalizes this residual only where $|\bar{v}| \geq 1$; the grey band is omitted from ratio claims. The residual is a group-theoretic matched counterfactual, not a causal label or an energetic definition of propulsion.

Population medians and upper quantiles are more informative here than extrema alone. For the anti-phase-descended family, the median absolute log-translation residual is 1.28793×10^{-3} . Among its 2,220 guarded cycles, the median normalized fraction is 0.007499%, the 95th percentile is 0.02807%, and every value is below 0.1%. For the in-phase-descended family, the median absolute residual is 2.50858×10^{-3} . Among its 1,350 guarded cycles, the normalized median is 0.003776%, the 95th percentile is 0.5921%, 98.37% lie below 1%, and the maximum is 1.497%.

At the nearest recorded pair to $\bar{v} = -8$, the anti-phase-descended cycle has $\bar{v} = -7.99366$ and $r_{\text{inc}} = 0.02414\%$; the in-phase-descended cycle has $\bar{v} = -8.06400$ and $r_{\text{inc}} = 0.2906\%$. This comparison quantifies different path modulation at nearly matched mean speed. It does not use exactly the same in-phase sample as the common-duration illustration, which was selected for visual alignment. It also does not establish which internal coordinate caused the displacement or how much energetic propulsion should be assigned to it.

Away from the near-zero guard, the matched straight term therefore dominates translation for most cycles, especially in the anti-phase-descended family, with the quantified in-phase-descended

tail above. No normalized dominance statement is made for $|\bar{v}| < 1$; absolute residuals must be read there. Repeating the reconstruction with every second represented sample changed the incremental log translation by less than 0.1% relative in the decimation check.

5.8 What transfers from the scalar construction

The carrier-closed auxiliary field, Schur-complement storage, local modal seeds, closed-to-open lift, full-field correction, and pose reconstruction all transfer from the scalar construction as useful organizing devices. The scalar exchange–return equivalence does not. In 3SEG, one storage value leaves a set of possible endpoints, so the opened state must return as a vector. Modal metrics classify the already-returned cycle, and exchange records the associated transverse-storage balance.

The numerical result is coexistence at one parameter vector: a faster, larger-excursion anti-phase-descended branch chain and a slower, smaller-excursion in-phase-descended branch chain, both consisting of returned moving cycles of the same unforced field. The computed data support the opened return, modal, field-consistency, exchange, and pose quantities described above; the missing family-wide closed-support genealogy has already been separated from these cycle-level results.

Vector return separates the two modal families

The scalar theorem of [Chapter 4](#) stops at one effective transverse degree of freedom because a regular storage level in the multi-transverse section still contains a surface of possible endpoints. Six-component return therefore qualifies each relative-periodic solution; full-cycle modal metrics then classify it as in-phase- or anti-phase-descended; zero cycle-mean exchange follows as a balance identity once that state has returned.

The reconstructed pose establishes moving branch chains, while the matched-straight SE(2) comparison shows that their translation is usually dominated by the same-speed straight contribution outside the declared near-zero guard. The remaining group residual measures path modulation, not a unique causal or energetic contribution of internal oscillation.

These conservative cycles repeat without maintained input only when the ideal model is initialized exactly on them; no attraction has been established. [Chapter 6](#) next changes the energy equation by adding dissipation and periodic input, so exact work balance replaces conservative energy preservation.

Part III

Dissipation, resonantly maintained locomotion, and local regulation

Chapter 6

From passive families to resonantly maintained locomotion

The conservative families of [Chapters 3 to 5](#) persist in their ideal models only because no mechanical energy is lost. Once dissipation is introduced, the effective transverse oscillator and the propulsive channel may still exchange power, but their total stored energy decreases. A repeatable motion must then receive, over every cycle, exactly the work removed by the losses. This changes the natural question from *which unforced motions does the mechanism admit?* to *which responses can be maintained while the mechanism, rather than a prescribed joint choreography, still determines the finite-amplitude waveform?*

The input used here is deliberately modest: a sinusoidal torque is applied at the internal joint. It supplies energy and a clock, but it does not prescribe the joint trajectory, its harmonics, the carrier motion, or the accumulated pose. These quantities are determined by a synchronized solution of the full forced–damped equations. Nonlinear frequency-response continuation, cycle-work balance, quadrature diagnostics, folds, and coexisting responses provide the established tools for finding and interpreting such solutions ([Cenedese and Haller, 2020a,b](#); [Peeters et al., 2011](#); [Kuether et al., 2015](#); [Scheel et al., 2018](#)).

Two constructions must be kept separate throughout the chapter. The first adds loss and actuation to one conservative 2SEG cycle without changing its mechanical core; it provides a local mechanics-preserving bridge. The second is a much broader forced-response dataset with the same two-segment kinematic architecture but a different parameter vector and elastic law. It constructs a zero set of responses synchronized one-to-one with the forcing clock. This set $\mathcal{S}_R$, its transported subset $\mathcal{M}_R$, and the regular local components of $\mathcal{M}_R$ must be distinguished: only the last are called *Resonant Locomotion Manifold* (RLM) charts. Phase, waveform, speed, and pose are observable projections of these solutions, not their definition.

6.1 From energy loss to synchronized return

The scalar forced oscillator makes the energetic transition explicit:

$$m\ddot{x} + c\dot{x} + kx + \alpha x^3 = F \cos(\Omega t). \quad (6.1)$$

With $F = 0$ and $c > 0$, the mechanical energy decreases as $\dot{E} = -c\dot{x}^2$ and a non-equilibrium oscillation decays. With forcing,

$$\dot{E} = F \cos(\Omega t) \dot{x} - c\dot{x}^2 = P_{\text{in}} - P_{\text{diss}}. \quad (6.2)$$




Figure 6.1: The one-coordinate reference. In the lightly damped linear case, the response peak lies near a quarter-period displacement lag. At finite amplitude, a response projection can bend and become multivalued. These panels teach the diagnostic and the fold; they are not predictions for the locomotor models below.

If a response returns after the excitation period $T_{\text{exc}} = 2\pi/\Omega$, then its stored energy also returns and

$$\int_0^{T_{\text{exc}}} P_{\text{in}} dt = \int_0^{T_{\text{exc}}} P_{\text{diss}} dt. \quad (6.3)$$

This equality is necessary for a periodic response. It does not determine the response waveform, phase, existence, or stability.

The same accounting applies to the locomotor. Ideal constraint reactions still perform no total work: they restrict the admissible velocities and structure the coupling through which the internal oscillator and propulsive channel redistribute energy, but they do not replenish the energy removed by damping. With generalized loss and actuation forces,

$$\dot{E} = \dot{q}^T Q_{\text{loss}} + \dot{q}^T Q_{\text{act}} = -P_{\text{diss}} + P_{\text{in}}. \quad (6.4)$$

Internal exchange cancels in this total balance. It remains mechanically important because it determines how joint work is distributed through the admissible carrier-shape motion.

Sampling at one fixed phase of the input gives the stroboscopic map $P_{T_{\text{exc}}}$. A period-matched response is a fixed point,

$$P_{T_{\text{exc}}}(z_0) = z_0. \quad (6.5)$$

The forcing clock removes the free choice of time origin present on an autonomous orbit. Equation (6.5) states return at the same forcing phase; it does not state that the fixed point attracts nearby initial conditions.

Figure 6.1 recalls the simplest frequency-response picture before it is applied to the locomotor.

The linear panel explains why quadrature is useful: when displacement lags force by a quarter period, its velocity is in phase with the force and receives positive mean power. The nonlinear panel explains why that intuition must be checked rather than assumed. Several responses may coexist at the same input, and one phase value need not identify the same observable maximum on every part of a folded response set.

6.2 A mechanics-preserving bridge for one conservative cycle

The first calculation retains exactly the geometry, inertia, nonholonomic constraint, and quadratic-quartic potential of the conservative 2SEG model in Chapter 3. Table 6.1 shows the entire mechanical change: the projected conservative field is kept, and only the nonconservative generalized-force term is augmented.

Table 6.1: The mechanics-preserving transition. Geometry, ideal constraint, admissible-velocity basis, inertia, and potential are identical in all three columns; only the generalized nonconservative force changes.

	Conservative	Damped, unforced	Damped and forced
Generalized nonconservative force	$Q = 0$	$Q = Q_{\text{loss}}$	$Q = Q_{\text{loss}} + Q_{\text{act}}$
Energy implication	$\dot{E} = 0$	$\dot{E} = -P_{\text{diss}}$	$\dot{E} = -P_{\text{diss}} + P_{\text{in}}$

Starting from (3.10), the computed bridge ramps isotropic viscous loss and synchronized joint torque through the homotopy

$$M(\delta)\dot{\eta} + c(\delta, \eta) + [0 \quad 0 \quad U'(\delta)]^\top + \lambda c_\star \eta = B_\delta \lambda A_\star \sin\left(\frac{2\pi t}{T_\lambda} + \phi_\lambda\right), \quad B_\delta = [0 \quad 0 \quad 1]^\top. \quad (6.6)$$

Here $c_\star = 10^{-3}$ and $A_\star = 0.0371395$. Thus $\lambda = 0$ is the original unforced conservative field. For $\lambda > 0$, damping and forcing are ramped together, while the response period T_λ and input phase ϕ_λ are corrected with the orbit. The time origin is fixed by the oriented section $\delta(0) = 0$, $\sigma(0) > 0$.

The anchor is one numerically refined conservative cycle (stored identifier RG327) at

$$\bar{v} = 10.43581096, \quad T_0 = 0.67499472, \quad \|z(T_0) - z(0)\|_\infty = 2.22 \times 10^{-14}. \quad (6.7)$$

At each positive λ , four reduced-return rows and the same fixed mean speed are imposed. Nineteen prescribed values connect $\lambda = 0$ to $\lambda = 1$. The maximum strict return residual is 7.11×10^{-10} and the maximum mean-speed error is 1.08×10^{-12} . At the first positive point, $\lambda = 10^{-4}$, the normalized waveform distance from the conservative anchor is 4.66×10^{-5} ; the period and joint excursion differ by only $-2.07 \times 10^{-3}\%$ and $2.42 \times 10^{-3}\%$, respectively.

Figure 6.2 shows the deformation of the orbit, its cycle-work balance, and the matched-straight pose comparison along the same branch.

At $\lambda = 1$, the period is 0.60392282, a 10.529% decrease from the anchor, and the joint peak-to-peak excursion has increased by 13.724%. The input work and dissipated work are, respectively, 0.104563190334 and 0.104563190336; their relative difference is 2.60×10^{-11} . Across the computed branch, the mechanical energy identity has absolute error below 4.12×10^{-13} .

This calculation supplies a local synchronized connection from one conservative cycle under the same mechanical core. It does not extend that connection to the whole conservative family. The physical input phase becomes unidentifiable as $\lambda \rightarrow 0$ because the input vanishes. The homotopy equations and numerical checks are given in Chapter E.

6.3 A separate parameter-distinct forced model

The broad response computation answers a different question: how does a related two-segment locomotor organize finite-amplitude responses over a broad forcing domain? The model was introduced in an earlier forced-response study (Rajaomarasata et al., 2025b); the scientific object used here is the declared model and its computed response set.

The configuration and reduced state are

$$q_f = (g, \delta), \quad g = (x, y, \theta) \in \text{SE}(2), \quad z_f = (\delta, v, \omega_b, \sigma)^\top, \quad (6.8)$$

with carrier-forward velocity v , body-yaw rate $\omega_b = \dot{\theta}$, and joint rate $\sigma = \dot{\delta}$. The same wheel-like row removes lateral carrier velocity.

Figure 6.3 makes the shared kinematic architecture and the internal torque channel explicit.




Figure 6.2: One mechanics-preserving conservative-to-forced bridge at fixed mean speed. (a) Period and joint excursion vary continuously from the conservative anchor. (b) Phase-aligned joint waveforms approach the anchor as $\lambda \rightarrow 0$. (c) Input work and viscous dissipation coincide on the returned branch. (d) The translational component of $\log(\Delta g_{inc})$ is normalized by the matched straight translation. This last quantity is a reference-dependent pose residual, not an additive or causal attribution of propulsion.

Its projected equations use the verified dissipative sign convention

$$M_f(\delta)\dot{\eta}_f + F_f(\delta, \eta_f) + D_f\eta_f = B_\delta\tau_\delta(t), \quad \eta_f = [v \quad \omega_b \quad \sigma]^\top, \quad D_f = c_d I_3, \quad c_d = 10^{-3}. \quad (6.9)$$

The nondimensional potential is $V_f(\delta) = \delta^2$, so its restoring torque is -2δ . The design vector is $\varepsilon_f = 0.73170844$, $\gamma_f = 0.73680149$. Exact coefficients, scaling, and sign checks are given in Chapter E.

Table 6.2 prevents the local bridge and the broad response computation from being read as one continuation.




Figure 6.3: Architecture of the parameter-distinct forced model. Its kinematic pattern matches 2SEG, but its design vector and elastic law do not match the conservative reference or the bridge in Section 6.2.

Table 6.2: The local bridge and the broad forced-response dataset answer separate questions.

Feature	Same-mechanics bridge	Broad forced-response dataset
Kinematics	Same 2SEG constraint and reconstruction	Same kinematic type
Elastic and inertial design	Exactly the Chapter 3 design	Different design vector and quadratic potential
Forced object	One branch from one conservative anchor at fixed $\bar{v}$	Broad independently computed response set
Permitted inference	Local numerical mechanics-preserving connection	Synchronized-response zero set; local RLM chart only where transport, rank, and embedding conditions hold

The imposed torque and its mechanical power are

$$\tau_{\delta}(t) = A_{\text{exc}} \sin(\Omega t + \varphi_{\text{in}}), \quad \Omega = \frac{2\pi}{T_{\text{exc}}}, \quad P_{\text{in}} = \tau_{\delta} \sigma. \quad (6.10)$$

All reported amplitudes are nondimensional. The stored plotting coordinate is $10^3 A_{\text{exc}}$; it is an N mm value only under the optional scaling $k_0 = 1$ N m. Viscous dissipation is

$$P_{\text{diss}} = c_d (v^2 + \omega_b^2 + \sigma^2). \quad (6.11)$$

For any exact periodic solution of (6.9), these two powers satisfy (6.3). This is an analytical consequence of periodicity, not a second independent response condition.

The pose remains open and is reconstructed from

$$\dot{x} = v \cos \theta, \quad \dot{y} = v \sin \theta, \quad \dot{\theta} = \omega_b. \quad (6.12)$$

Thus a forced locomotor response can close in z_f while accumulating a pose increment $\Delta g = g(0)^{-1} g(T_{\text{exc}})$.

6.4 Synchronized-response sheets and their transported subset

At a fixed input phase, a synchronized response solves

$$\mathcal{R}_{\text{per}}(z_0; A_{\text{exc}}, T_{\text{exc}}) := \Phi_{T_{\text{exc}}}^{A_{\text{exc}}, T_{\text{exc}}}(z_0) - z_0 = 0. \quad (6.13)$$

Here the forcing clock fixes the time origin. An autonomous forcing lift and a scalar numerical phase gauge provide an equivalent periodic boundary-value formulation; that numerical gauge must not be confused with the measured output–input phase introduced below. The implementation is detailed in [Chapter E](#).

For the fixed forced model, define the synchronized-response zero set and its transported subset

$$\begin{aligned} \mathcal{S}_R &= \{(z_0, A, T) \in \mathbb{R}^4 \times \mathbb{R}_+^2 : \mathcal{R}_{\text{per}}(z_0; A, T) = 0\}, \\ \mathcal{M}_R &= \{s \in \mathcal{S}_R : \Delta g(s) \neq e_G\}, \quad G = \text{SE}(2). \end{aligned} \quad (6.14)$$

At a point of $\mathcal{S}_R$ where the Jacobian of the four independent return equations has rank four with respect to (z_0, A, T) , the implicit-function theorem gives a local two-dimensional solution sheet in orbit–forcing space. The nonidentity-pose condition is open, so the same local dimension holds in $\mathcal{M}_R$. A local component of $\mathcal{M}_R$ is called an RLM chart only where this rank condition holds and the solution sheet is embedded. Restoring the forcing phase associates a locally three-dimensional swept set in extended state–clock space; that sweep is not an additional coordinate of $\mathcal{S}_R$ or $\mathcal{M}_R$.

and does not alter the two-dimensional count above. A transverse fixed-amplitude slice gives a local frequency-response curve.

The notation is deliberately asymmetric. The zero set $\mathcal{S}_R$ contains synchronized responses whether or not they translate, whereas $\mathcal{M}_R$ imposes transport without presuming regularity. Population-wide integration of the recorded (v, ω_b) histories gives a nonidentity pose increment for all 13,782 continued responses; the reconstruction and integration comparison are reported in [Section E.5](#). An independent exact-field propagation check also covers every one of the 2,756,400 stored intervals, closure included, with the exact declared field; its largest scaled coordinate defect is 1.716×10^{-5} . At the reported numerical resolution, synchronized return and pose reconstruction therefore place every recorded response in $\mathcal{M}_R$. The definition still leaves open synchronized solutions with identity pose increment elsewhere in $\mathcal{S}_R$.

Here *resonant* identifies how the component is constructed: continuation starts near the small-amplitude natural period and follows the primary 1:1 response component through its nonlinear response ridge. The unit winding verifies synchronization, not resonance by itself. Nor must every member be located at a speed maximum or at quadrature. Complete collocation Jacobians are unavailable for a branch-wide regularity test, so the RLM name applies to local regular charts rather than to $\mathcal{M}_R$ globally through possible singular points.

The response population was assembled by continuation in excitation period at fixed amplitude, followed by continuation in amplitude from selected responses. Pseudo-arclength correction follows the zero set of (6.13) through folds, as introduced in [Chapter 2](#). It produces an orbit-valued solution family, not the asymptotic trace generated by one chosen initial condition.

Measured quantities are coordinates on, or projections of, the underlying response set. Using total harmonic distortion (THD) for deformation waveform content, define

$$\begin{aligned} \mathcal{O} : \mathcal{S}_R &\longrightarrow \mathbb{R}_+^2 \times \mathbb{R}^4 \times \text{SE}(2), \\ (z_0, A, T) &\longmapsto (A, T, v_{\text{pk}}, \bar{v}, \phi_{\delta\tau}, \text{THD}_\delta, \Delta g), \end{aligned} \quad (6.15)$$

where v_{pk} is the sampled positive within-cycle maximum, $\bar{v} = T^{-1} \int_0^T v(t) dt$ is the cycle-mean body-frame longitudinal speed, $\phi_{\delta\tau}$ is the deformation phase relative to torque, and THD_δ is the deformation THD. The plotted set is an observable image $\mathcal{O}(\mathcal{S}_R)$; restricting it to transported responses gives $\mathcal{O}(\mathcal{M}_R)$. Either image may fold, self-intersect, or lose rank even where the solution sheet is smooth. Any such local, parameter-distinct RLM chart is therefore neither the conservative Natural Locomotion Manifold nor the mechanics-preserving bridge above.

6.4.1 Phase measured from deformation and torque

The offline phase uses joint deformation as output and applied torque as input. Each adaptive-mesh cycle is first interpolated periodically onto a uniform grid with the duplicated endpoint removed. Centered analytic signals are then used to form the circular mean of their pointwise phase difference. The convention is

$$\phi_{\delta\tau} = \text{phase}_H(\delta) - \text{phase}_H(\tau_\delta) \quad \text{wrapped to } [-\pi, \pi), \quad \phi_{\delta\tau}^Q = -\frac{\pi}{2}. \quad (6.16)$$

Here phase_H is shorthand for the specified Hilbert analytic-signal statistic, not a pointwise phase on an arbitrary transient. The resampling, circular-mean formula, and harmonic extraction are specified in [Chapter E](#). [Chapter 7](#) uses a different online fundamental estimator with the same signal order and target.

6.5 Speed and period across forcing-amplitude levels

At fixed amplitude, the positive peak speed varies with excitation period, and the response curves bend enough that the same input coordinates can correspond to several period-matched responses.




Figure 6.4: Positive peak-forward-speed responses for the lowest 15 forcing-amplitude levels, $0.00204965 \leq A_{\text{exc}} \leq 0.03551893$. Each curve is a continued set at one fixed nondimensional amplitude; the integer beside its peak orders the levels from smallest to largest, so the identification does not depend on color. Circles mark sampled speed maxima, crosses mark piecewise-linearly interpolated $-\pi/2$ phase crossings, and the dashed line is the exact linearized period. Color denotes the display coordinate $10^3 A_{\text{exc}}$.

Figure 6.4 displays this projection for the lowest 15 forcing-amplitude levels, $0.00204965 \leq A_{\text{exc}} \leq 0.03551893$.

The computation contains 13 782 continued cycles at 25 fixed amplitude levels over $A_{\text{exc}} = 0.00204965$ – 0.07357499 ; the complete list is given in Equation (E.15). The mechanical and forcing-clock states return at the recorded precision for every cycle. The seam-inclusive exact-field and cycle-work checks in Chapter E independently qualify the stored traces; the original continuation Jacobians remain unavailable for a branch-wide rank test.

At the straight rest state, eliminating yaw gives the effective shape inertia $M_{\text{lin}} = 0.02582240489$. Since $V_f = \delta^2$, the linear period is

$$T_{\text{lin}} = 2\pi \sqrt{\frac{M_{\text{lin}}}{2}} = 0.7139424640. \quad (6.17)$$

The observed peak-period shift is finite and non-monotone: 3.7940% at $A_{\text{exc}} = 0.03551893$, 4.4691% at $A_{\text{exc}} = 0.03969468$, and 3.8114% at $A_{\text{exc}} = 0.07357499$. These values establish the response-level departure from the linear reference without assigning it to one term of the equations.

A fold means that a vertical line in Figure 6.4 can meet one fixed-amplitude curve more than once. These intersections are distinct periodic solutions of the boundary-value problem at the same forcing coordinates. They need not share waveform, mean speed, phase, or pose increment. Their coexistence does not imply that every solution is attracting or reachable from the same initial condition.

6.6 Exact cycle work and a local quadrature marker

The exact and empirical statements must remain separate. For every exact periodic response of the forced model, full-cycle work balance is

$$\int_0^T \tau_\delta \sigma \, dt = \int_0^T c_d(v^2 + \omega_b^2 + \sigma^2) \, dt. \quad (6.18)$$

It includes the complete waveform and all three dissipative pseudo-velocity channels. Although the actuator is collocated with the shape rate through $P_{\text{in}} = \tau_\delta \sigma$, the admissible inertial coupling redistributes that work between the effective oscillator and the propulsive channel. The ideal constraint structures this transfer without contributing power to the total balance.

Quadrature has a narrower role. If the torque and deformation fundamentals are

$$\tau_1(t) = \hat{\tau} \cos(\Omega t), \quad \delta_1(t) = \hat{\delta} \cos(\Omega t - \pi/2), \quad (6.19)$$

then $\dot{\delta}_1$ is in phase with the torque and

$$\frac{1}{T} \int_0^T \tau_1 \dot{\delta}_1 \, dt = \frac{1}{2} \Omega \hat{\tau} \hat{\delta} > 0. \quad (6.20)$$

In the declared sign convention, positive torque aligned with positive shape velocity injects work and counters the negative viscous contribution. This fundamental-harmonic calculation motivates the sign and the use of phase. It is neither a pointwise condition on the nonlinear waveform nor an equivalent form of the complete work identity in (6.18); the phase statistic itself is defined in (6.16).

Let $\mathcal{A}_{\text{low}}$ denote the lowest 15 amplitude levels, $0.00204965 \leq A_{\text{exc}} \leq 0.03551893$. At each level, the comparison uses the sampled maximum (T_r, v_r) with the piecewise-linearly interpolated $-\pi/2$ crossing (T_Q, v_Q) . Every level in $\mathcal{A}_{\text{low}}$ contains such a crossing, and

$$\max_{A_{\text{exc}} \in \mathcal{A}_{\text{low}}} 100 \frac{|v_Q - v_r|}{v_r} = 0.0635625\% < 0.064\%, \quad \max_{A_{\text{exc}} \in \mathcal{A}_{\text{low}}} 100 \frac{|T_Q - T_r|}{T_r} = 0.0289515\%. \quad (6.21)$$

The largest phase distance from a sampled speed maximum to quadrature on this range is 0.0452127 rad. These bounds concern one estimator, one sampled domain, and the positive peak-speed observable.

The behavior changes at the next five levels, $0.03969468 \leq A_{\text{exc}} \leq 0.05655617$. They contain a sampled quadrature crossing on a branch whose peak speed is approximately 59–72% below the amplitude-level maximum. At the five upper levels, $0.06052866 \leq A_{\text{exc}} \leq 0.07357499$, the sampled curves do not cross $-\pi/2$.

Figure 6.5 places the speed maximum, deformation harmonics, and phase statistic on the same amplitude axis.

The same figure also shows why an imposed sinusoidal torque does not prescribe a sinusoidal locomotor state. The joint waveform acquires harmonics, and the relation among shape, carrier speed, and yaw changes over the response set. Quadrature is therefore a useful local coordinate on $\mathcal{A}_{\text{low}}$, not a global definition of resonance or locomotor performance.




Figure 6.5: Observables at the sampled speed maximum of all 25 forcing-amplitude levels. Peak speed grows and plateaus (left), deformation harmonic content changes (center), and the output-minus-input Hilbert phase remains near $-\pi/2$ only on the lower-amplitude segment (right). The horizontal coordinate is the nondimensional display value $10^3 A_{\text{exc}}$.

6.7 Pose reconstruction separates locomotion from oscillation

Peak velocity alone does not establish net motion: an oscillatory velocity can integrate to zero. The pose equation (6.12) supplies the required test. Population-wide reconstruction verifies the nonidentity-pose condition for all recorded traces. Combined with their population-wide dynamic-return check, it places them numerically in $\mathcal{M}_R$. Regular embedded local components of this set are the RLM charts. The three representative responses in Figure 6.6 show the corresponding paths at low, intermediate, and high forcing amplitudes.




Figure 6.6: Pose reconstructions for peak-speed responses at $A_{\text{exc}} = 0.00204965$, 0.02763993 , and 0.07357499 . Artificial vertical offsets separate the paths. Each displayed reduced response returns over one forcing period while the reconstructed carrier advances.

The positive-peak-speed ridge is not automatically an optimum of mean translation, input work per distance, actuator effort, or stability. Nor does phase alone locate a response on a folded map: two responses can have similar phase and different mean speeds. This ambiguity motivates the two coordinates used in [Chapter 7](#): phase adjusts forcing frequency on the fast loop, while mean forward speed adjusts amplitude on the slower loop.

From resonant maintenance to regulation

Dissipation changes the repeatable object because stored energy returns only when cycle input work replaces cycle loss. The mechanics-preserving homotopy shows this transition for one conservative 2SEG cycle. The parameter-distinct continuation is designed to sample $\mathcal{S}_R$, and pose reconstruction selects its transported subset $\mathcal{M}_R$ after independent local-flow qualification; regular embedded local components of the latter are RLM charts. In both cases, synchronized return determines the nonlinear response; the sinusoidal input supplies work and timing rather than a complete choreography.

Two statements then organize the passage to control. Work balance is exact for every exact periodic response. Output-minus-input quadrature is a local empirical marker of the positive-peak-speed ridge over $0.00204965 \leq A_{\text{exc}} \leq 0.03551893$; at larger amplitudes it either belongs to a slower branch or is absent from the sampled curves. Because folds allow similar phase values on dynamically distinct responses, phase alone cannot locate a point on the response sheet. The next chapter pairs this timing coordinate with mean forward speed on the finite tested segment.

Chapter 7

Local phase–speed regulation of resonant responses

The synchronized-response sheets of [Chapter 6](#) organize the periodic solutions; their regular transported portions admit conditional local Resonant Locomotion Manifold (RLM) charts. Neither object is itself a feedback law, and measured phase and speed do not decide membership in a particular stored response. On the lower-amplitude sheet, two measured quantities supply complementary local information. The output–input phase records whether joint work is timed near the observed quadrature relation; mean body-frame longitudinal speed distinguishes operating points along the useful sampled rising relation.

This division leads directly to the control architecture. A fast frequency adjustment regulates phase, a slower amplitude adjustment moves mean speed, and guards restrict the motion to the finite region in which those coordinates were tested. A phase-locked loop (PLL) provides the first layer, following established uses of phase feedback for nonlinear response tracking and control-based continuation ([Peter et al., 2016](#); [Peter and Leine, 2017](#); [Denis et al., 2018](#); [Renson et al., 2016](#); [Barton, 2017](#); [Abeloos et al., 2022](#); [Tang et al., 2020](#)). The second layer is specifically locomotor: phase alone is not a branch coordinate when several responses share nearly the same timing relation.

The objective is also adjacent to orbital stabilization of underactuated mechanical systems, where feedback is designed transverse to a prescribed periodic orbit rather than to a time-indexed trajectory ([Canudas-de Wit et al., 2002](#)). Hybrid zero dynamics applies a related orbit-and-invariant-subsystem viewpoint to impact-driven bipedal locomotion ([Westervelt et al., 2003](#)). Those methods supply an important boundary, not a result claimed here: the present continuously constrained model uses neither a virtual constraint nor an impact/reset map, and the two observable loops below do not prove orbital stability.

Whole-body robot controllers address a broader coordination problem by combining equality and inequality constraints in strict task hierarchies, hierarchical quadratic programs, and rigid-contact dynamics ([Mansard et al., 2009](#); [Escande et al., 2014](#); [Saab et al., 2013](#)). The architecture here is deliberately narrower: two response observables selected by the locomotor mechanics are regulated locally, and guards delimit the tested response segment; no whole-body task stack or contact-force optimizer is prescribed.

The resulting controller regulates phase and mean speed without querying the offline response family. Its operating range is the rising response segment identified in [Chapter 6](#); the response family is used after each run only to interpret the motion that was reached.

7.1 Why phase is an ambiguous coordinate

Phase measures timing, not the complete response. Two periodic motions can place the fundamental deformation nearly one quarter period behind the torque while differing in deformation amplitude, basin, harmonics, and mean longitudinal speed. In a folded response set, the same phase value may therefore occur on several dynamically distinct responses.

The forced-model data expose this ambiguity directly. An amplitude-staircase experiment attempted to reach the stored response at $A_{\text{exc}} = 0.0176277$, whose mean speed is 6.2112, while retaining the PLL. Fifteen trials used several low-response starts and target-phase settings. None reached that speed neighborhood: all 15 settled on a low-speed response, even though their largest root-mean-square (RMS) phase error was only 0.07138 rad. A small phase error can consequently indicate lock to a response without indicating lock to the intended response.

Let Γ denote a settled forced response. The two observables used here are

$$\Gamma \mapsto (\phi(\Gamma), \bar{v}(\Gamma)), \quad \bar{v}(\Gamma) = \frac{1}{T} \int_0^T v(t) \, dt. \quad (7.1)$$

Here $v(t)$ is the longitudinal velocity resolved in the body frame. Hence $\bar{v}$ is a period-averaged local control coordinate, not, when the heading changes, the norm of the reconstructed inertial translation divided by T . Net transport is inferred only for responses whose full pose has been reconstructed. The first component of Equation (7.1) compares response timing with the applied torque; the second locates a response on the tested speed segment. Order the stored responses by increasing forcing amplitude; an *amplitude level* below means one of these stored values, not an online index used by the controller. On the first fourteen sampled levels, the pairs

$$(A_{\text{exc}}, \bar{v}) : (0.0020496, 0.021457) \longrightarrow (0.0314870, 10.876395),$$

increase, and every adjacent sampled slope is positive. Mean speed is therefore a usable local coordinate on this sampled segment: it separates responses whose phases are almost indistinguishable. The next amplitude level, $(A_{\text{exc}}, \bar{v}) = (0.0355189, 12.364852)$, supplies the numerical high-speed endpoint. The present control tests do not extend beyond this declared observable interval.

The two measurements now answer different questions. Phase asks whether the current timing is near the local quadrature marker. Mean speed asks where the settled motion lies on the useful locomotor segment.

7.2 Online phase and the locally tested frequency sign

The controller observes the internal deformation δ relative to the applied joint torque

$$\tau_\delta(t) = A_{\text{exc}} \sin \varphi_{\text{exc}}(t), \quad \dot{\varphi}_{\text{exc}} = \Omega. \quad (7.2)$$

On a rolling window of one estimated forcing period, a three-harmonic least-squares fit absorbs waveform distortion. Only its fundamental phases enter feedback:

$$\phi = \text{wrap}_{[-\pi, \pi)}(\phi_\delta - \phi_\tau), \quad e_\phi = \text{wrap}_{[-\pi, \pi)}\left(\phi + \frac{\pi}{2}\right). \quad (7.3)$$

The order is output minus input. A deformation fundamental delayed by one quarter period has $\phi = -\pi/2$ and hence $e_\phi = 0$, as illustrated in Figure 7.1. Estimator algebra, window definitions, and conditioning diagnostics are given in Chapter F.




Figure 7.1: Online phase convention. The applied torque provides the excitation clock, and a rolling harmonic fit extracts the fundamental deformation phase. In the output-minus-input convention, the displayed quarter-period delay is $-\pi/2$. Higher fitted harmonics absorb waveform distortion; they are not additional feedback coordinates.

This online observable is not numerically identical to the offline Hilbert phase used to build the observable response projection in [Section 6.4.1](#). Both use the same output-minus-input order and quadrature target, but transients, windows, and waveform distortion affect them differently. The feedback sign was therefore tested on the online observable rather than inferred from the offline plot.

A condensed proportional–integral (PI) form of the phase controller is

$$\dot{I}_\phi = e_\phi, \quad \Omega_{\text{raw}} = \Omega_{\text{free}} + s(K_P e_\phi + K_I I_\phi), \quad \Omega = \text{sat}_{[5,11]}(\Omega_{\text{raw}}). \quad (7.4)$$

Here s is fixed once [Equation \(7.3\)](#) has been chosen. Around the amplitude level $A_{\text{exc}} = 0.0176277$, with $\Omega_{\text{exc}} = 8.7439$, a positive e_ϕ must increase Ω , hence $s = +1$. With that sign, initial detunings of -0.5% , $+0.5\%$, and -3% ended with phase-error RMS below 4.2×10^{-3} rad and mean frequency close to the reference. With the opposite sign and the same -0.5% detuning, the phase-error RMS remained 1.2452 rad and the mean frequency fell to 0.874 of the reference. This calibrates the feedback sign in that local neighborhood; the tested detunings are not a map of its acquisition basin. The exact sampled PI recurrence, including estimator startup and frequency saturation, is given in [Section F.2](#).

[Figure 7.2](#) compares the correct and reversed feedback signs under the same local perturbation.




Figure 7.2: Local sign experiment at $A_{\text{exc}} = 0.0176277$. Correct-sign feedback reduces the phase error for the three tested frequency detunings and returns the mean frequency near its reference. Reversing the sign under the same -0.5% detuning moves away from that neighborhood.

7.3 Fast phase loop, slow speed loop, and guards

The two-loop architecture follows from the two-coordinate problem:

$$\underbrace{e_\phi \mapsto \Omega}_{\text{fast phase loop}}, \quad \underbrace{e_v = \bar{v}_{\text{des}} - \bar{v} \mapsto A_{\text{exc}}}_{\text{slow speed loop}}. \quad (7.5)$$

Changing Ω corrects the response–forcing timing near the current periodic motion. Changing A_{exc} slowly changes which settled response is maintained. The second loop is therefore not an instantaneous velocity servo: it estimates mean speed over a rolling three-period window and moves the operating point while the faster loop preserves phase lock.

The phase error also gates the amplitude rate. For the nominal campaign,

$$g_\phi(e_\phi) = \begin{cases} 1, & |e_\phi| \leq 0.20, \\ \frac{0.55 - |e_\phi|}{0.35}, & 0.20 < |e_\phi| < 0.55, \\ 0, & |e_\phi| \geq 0.55, \end{cases} \quad (7.6)$$

and the slow law is

$$e_v = \bar{v}_{\text{des}} - \bar{v}, \quad \dot{A}_{\text{exc}} = \text{sat}_{[-r_A, r_A]} \left(g_\phi K_v e_v - \mathbf{1}_{\{g_\phi=0\}} \rho (A_{\text{exc}} - A_{\min}) \right), \quad A_{\text{exc}} \in [A_{\min}, A_{\max}]. \quad (7.7)$$

The tested values are $K_P = 0.04$, $K_I = 0.08$, $K_v = 2.5 \times 10^{-5}$, $r_A = 2.5 \times 10^{-4}$, $\rho = 10^{-4}$, and $[A_{\min}, A_{\max}] = [10^{-4}, 8 \times 10^{-2}]$. Before a valid phase estimate exists, the amplitude update is held. The sampled recurrences and the distinction between rolling estimator windows and end-of-stage evaluation windows are stated in [Sections F.2 and F.4](#).

[Figure 7.3](#) shows the information exchanged by the two loops and the stage-transition logic.




Figure 7.3: Observable-level architecture. Frequency is the fast phase actuator; amplitude is the slow locomotor-coordinate actuator. Guards inspect completed stages and decide which terminal states are retained. Ordered forcing-amplitude levels are used only for post-run interpretation.

A component ablation at desired mean speed $\bar{v}_{\text{des}} = 6$, reported in [Table 7.1](#), makes the separation concrete.

Table 7.1: Component ablation at the speed-6 command.

Active layers	Achieved $\bar{v}$	Phase-error RMS	Observed role
Phase and speed loops	5.594	8.14×10^{-3}	Moves upward while retaining phase lock
Fixed-amplitude PLL	1.798	7.07×10^{-4}	Locks phase at the starting response
Speed loop without PLL	3.001	0.812	Moves speed but loses the phase marker

The phase loop can preserve the wrong locomotor operating point, while the speed loop alone can leave the local resonant marker.

After each finite-duration stage, the stage manager evaluates phase-error RMS, deformation peak-to-peak, mean yaw rate, backward speed progress, and, above the high-speed threshold, proximity to the command. A passing stage retains its terminal plant state, excitation phase, amplitude, and frequency. A failing stage restores those last accepted values and halves the proposed speed increment. Exact definitions, thresholds, durations, and step-update rules are collected in [Section F.4](#). The tests use measured trajectory quantities only; the stored forcing-amplitude levels are attached after a run for interpretation.

7.4 Regulating phase and speed across the rising interval

The nominal campaign used two starts: rest, supplied with the forcing coordinates $(A_{\text{exc}}, \Omega) = (0.0020496, 8.8006)$, and the stored periodic state at the same amplitude level. Each run accepted 21 stages and reached desired-speed command 12. The terminal mean speeds were 12.068039 and 12.067794, with terminal phase-error RMS 1.119×10^{-3} and 1.106×10^{-3} rad, respectively. The corresponding nondimensional final forcing amplitudes were 0.0364166 and 0.0364088; by mean speed, both endpoints are closest to the stored level at $A_{\text{exc}} = 0.0355189$. [Figure 7.4](#) shows the two stage sequences.




Figure 7.4: Nominal phase-speed progression from rest and from the stored weak-amplitude response at $A_{\text{exc}} = 0.0020496$. Both runs contain 21 accepted stages and terminate near mean speed 12.068. Once the accepted target reaches 8, the proposed increment is capped at 0.25. All phase-error RMS values remain below the acceptance threshold; the closest stored high-amplitude response has $A_{\text{exc}} = 0.0355189$ and is identified only after the run.

The mean-speed annotation is corroborated by a separate full-cycle comparison for the stored-response start. The last controlled forcing cycle and each of the 25 stored peak-response cycles were resampled at the same 512 forcing phases in the full reduced state $(\delta, v, \omega, \sigma)$. After componentwise scaling by the range of the 25 reference cycles, the closest complete profile was forcing-amplitude level 15, with normalized RMS distance 1.7543×10^{-2} ; amplitude level 14 was next at 6.5033×10^{-2} . The last two controlled cycles differed by 1.1275×10^{-4} , and the normalized one-cycle endpoint defect was 2.3416×10^{-4} . This a posteriori neighborhood test uses a complete cycle rather than phase and speed alone. It is not an exact return residual or a proof of membership in the full synchronized-response zero set; its metric and time-step check are specified in [Section F.5](#).

All 42 nominal stage trials were accepted, so those two runs do not exercise rejection. A separate stress harness deliberately bypassed the high-speed step cap after the accepted target 8.2935. The direct proposal to 12 was rejected at mean speed 8.9980. The harness then discarded that terminal state, explicitly reused a copy of the last accepted plant and forcing state, halved the proposed

increment, and accepted the resulting target 10.1468 at mean speed 8.7457. This tests the simulator, the guard, and one constructed reject–reuse–reduced-proposal sequence; it does not execute the rejection branch of the production stage manager or establish recovery after arbitrary failures. The two nominal endpoints remain the evidence for consistency across the named initializations.

7.5 Sampled reduced-order evidence

A separate calculation projects each sampled forced response onto $K_G = 5$ harmonics of its four reduced-state components, giving 44 coordinates, and linearizes the phase-only and two-loop closures under the tested gains. The resulting closed matrices are displayed in [Section F.6](#). The spectral abscissa is the largest real part among a matrix’s eigenvalues; a negative value means asymptotic stability of that finite linear model. This Galerkin order is unrelated to the online estimator order $K_{\text{est}} = 3$.

[Table 7.2](#) reports the resulting spectral-abscissa ranges on the sampled branch segment.

Table 7.2: Spectral-abscissa ranges of the sampled $K_G = 5$ Galerkin controller matrices.

Augmented model	Spectral-abscissa range	Negative samples	Domain
PLL frequency loop	$[-1.81848, -1.00100] \times 10^{-3}$	14/14	first 14 amplitude levels
PLL and slow speed loop	$[-8.93731, -2.95865] \times 10^{-4}$	14/14	first 14 amplitude levels

All 14 matrices, spanning $A_{\text{exc}} = 0.0020496$ to 0.0314870 , have negative spectral abscissa. The next level, $A_{\text{exc}} = 0.0355189$, is not one of these linearized samples. In particular, the controlled endpoints near $A_{\text{exc}} = 0.03641$ lie outside this finite-harmonic interval; successful terminal simulations are not covered by the spectral table. The matrix construction, the local unsaturated assumptions, and the projection-residual comparison are given in [Section F.6](#).

7.6 Control result and scope

Phase and speed solve different local problems. Frequency feedback regulates the response–forcing timing, slow amplitude adaptation moves the mean body-frame longitudinal speed, and the guards restrict accepted stages to those satisfying the declared observable thresholds. Their combination reaches the same high pre-transition neighborhood from the two named starts, while phase-only tests can remain locked to the wrong low-speed response.

The supported result is nominal phase–speed progression from two initializations, a complete-cycle neighborhood check at one endpoint, one constructed reject–reuse–reduced-proposal stress test, and sampled reduced-order evidence on the smaller declared amplitude interval. Repeating the final stored-response stage with half the integration step preserves its terminal observables and closest stored forcing-amplitude level. These finite tests do not establish a global acquisition basin, general recovery, arbitrary-target convergence, or stability of the exact nonlinear forced system. [Chapter 8](#) places this local regulation result in the complete conservative-to-forced argument.

Part IV

Synthesis, limits, and outlook

Chapter 8

Synthesis: accompanying intrinsic locomotor families

8.1 The scientific result

The thesis began with a change of design question. Instead of prescribing a complete joint choreography and asking the mechanism to reproduce it, one can first ask which motions are already compatible with its mechanics. In a fixed ideal conservative model, a natural-locomotion cycle is therefore a nontrivial unforced solution whose pose-reduced state returns exactly while its planar pose remains open and accumulates a declared nonidentity increment. A member initialized exactly on such a family persists in that ideal field because the family is invariant. This statement does not require, and does not imply, attraction or robustness.

Four successive insufficiencies determine the construction. A linear mode is only a small-amplitude seed: frequency, waveform, harmonics, and the phase relations between coordinates change at finite amplitude. Periodicity of the complete state is too restrictive for locomotion because it closes the pose and removes net rigid-body transport. An arbitrary loop in internal shape space is too permissive because it may reconstruct a displacement without solving the unforced equations with the prescribed timing. Finally, a conservative cycle is not maintained after positive losses are introduced; the repeatable object then becomes a synchronized forced response whose input work replaces the dissipated energy.

This sequence yields the central answer. Natural locomotion is neither a preferred sinusoid nor an arbitrary transporting loop. It is a dynamically compatible relative-periodic motion, organized in families and interpreted through the exchange between an effective transverse oscillator and a propulsive carrier channel. Under dissipation, the same principle becomes resonant accompaniment: inject the work needed by the response, then regulate phase and mean-speed observables over the interval where they have been tested.

8.2 One constrained-mechanics method, two locomotors

The two-segment locomotor (2SEG) and three-segment locomotor (3SEG) are derived by the same projected Lagrange–d’Alembert pattern. For a smooth ideal velocity constraint,

$$A(q)\dot{q} = 0, \quad \dot{q} = H(q)\eta, \quad \text{Im } H(q) = \ker A(q), \quad (8.1)$$

the admissible equations follow by projecting the Lagrange–d’Alembert equations onto the kernel basis,

$$H(q)^\top \left[\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q \right] = 0. \quad (8.2)$$

The symbol N used for the 3SEG basis has the same role as H in the 2SEG derivation. The reaction force disappears from the projected equations because it has zero power on every admissible velocity. The constraint nevertheless remains dynamically active: its distribution and the kinetic metric determine how the admissible carrier and joint velocities are coupled. The full derivations are given in Chapters B and D.

Projection and pose reconstruction are different operations. Projection produces the reduced mechanical field. After that field has been integrated, the group equation reconstructs (x, y, θ) and supplies the pose increment Δg . Leaving this pose outside the return residual is what allows a closed reduced cycle to describe locomotion rather than oscillation in place.

The pendulum analogies become precise at the level of transverse dynamics, not at the level of the complete mechanisms. The 2SEG has four configuration coordinates and three admissible velocity components; eliminating its two carrier components leaves one effective transverse oscillatory degree of freedom, analogous to the simple pendulum. The 3SEG has five configuration coordinates and four admissible velocity components; carrier elimination leaves two transverse oscillatory degrees of freedom, analogous to the double pendulum. Thus the simple/double-pendulum correspondence explains the scalar-to-vector transition without describing either locomotor as a globally one- or two-degree-of-freedom system.

The generality of this modelling pattern has a clear boundary. It applies to smooth, bilateral, ideal, continuously active, constant-rank velocity constraints. Impacts, unilateral contacts, changing contact modes, and stick-slip require additional hybrid or nonsmooth mechanics rather than a change of notation in (8.1) (Brogliato, 2016; Prasad et al., 2023).

8.3 Families, transport, and local manifold geometry

For a fixed conservative vector field, shooting and continuation solve an independent phase-fixed reduced-return problem

$$F_N(u, \lambda) = 0, \quad F_N : \mathbb{R}^{m+1} \longrightarrow \mathbb{R}^m. \quad (8.3)$$

Here m counts the independent return equations after the time-shift redundancy has been removed. If DF_N has rank m , the local zero set is a one-dimensional smooth branch. Restoring the phase along each nontrivial periodic orbit sweeps a locally two-dimensional invariant family, provided that the phase action is regular and the family is embedded. Pose reconstruction selects the open-pose subset $\Delta g \neq e_G$. In this precise, conditional sense, its regular embedded branches provide local charts of the Natural Locomotion Manifold (NLM). Neither a finite point cloud nor the image of a branch in a speed-energy plot establishes this regularity by itself.

This local language is intentionally stronger than using *manifold* as a metaphor and narrower than a global topology claim. Separate modal sectors, bifurcating branches, the 2SEG and 3SEG state spaces, and distinct parameter vectors are not merged into one smooth object. A fold of an observable projection may occur while the underlying solution curve remains regular; a branch point or rank loss requires a different local analysis.

Reduced return qualifies the dynamics, whereas pose reconstruction evaluates the locomotor output. All 379 exported 2SEG records have nonidentity reconstructed pose; independent exact-field and local-rank evidence covers 29 matched roots rather than the complete window. The returned in-phase- and anti-phase-descended 3SEG populations also accumulate finite pose increments on the displayed examples. Mean speed is a useful family label but not the naturality criterion. Moreover, the matched-straight SE(2) comparison shows that, away from nearly zero baseline displacement, straight relative-equilibrium translation is usually the dominant baseline in the tested families. The oscillatory motion modulates that moving baseline; the comparison does not assign a unique causal share of propulsion to the oscillation.

8.4 Energy exchange: exact scalar closure and its vector boundary

Carrier elimination in the 2SEG kinetic energy produces a completed-square decomposition

$$E = E_{\text{prop}} + E_{\perp}, \quad \dot{E}_{\perp} = P_{\text{ex}}, \quad \dot{E}_{\text{prop}} = -P_{\text{ex}}. \quad (8.4)$$

along the opened ideal field. Here $E_{\perp}$ is the storage of the effective transverse oscillator and E_{prop} is the complementary carrier storage. More explicitly, with $y = (v, \omega)^{\top}$ and the energy-minimizing carrier velocity y_* ,

$$E_{\text{prop}} = \frac{1}{2}(y - y_*)^{\top} M_{cc}(y - y_*), \quad y_* = -M_{cc}^{-1} m_{c\sigma} \sigma. \quad (8.5)$$

Thus E_{prop} is the energy of the defect $y - y_*$, not the complete carrier kinetic energy. On the closed leaf, $y = y_*$ is generally nonzero, and the Schur-complement storage $E_{\perp}$ already contains the carrier motion induced by the shape rate. The adjective *propulsive* therefore refers to the complementary channel in this declared carrier–shape partition; it is not a coordinate-free causal allocation of displacement. The ideal constraint supplies no energy: it structures the admissible coupling that transfers power between these two terms while their sum remains constant.

The mechanics construction has four distinct steps. Carrier elimination defines the reduced storage; the carrier-closed auxiliary field conserves that storage; a closed-to-open lift embeds a candidate in the opened variables; and restoring the original coupling gives the full field whose difference defines P_{ex} . Only after a reduced motion has been obtained is the pose reconstructed. This vocabulary prevents carrier reopening from being confused with group reconstruction.

On the positive oriented scalar section of the 2SEG, transverse storage is strictly increasing in the remaining crossing speed. Once the two carrier return rows close, zero integrated exchange is therefore equivalent to the last scalar return row. [Theorem 4.2](#) makes this necessity-and-sufficiency statement exact under its finite-speed, section, and injectivity hypotheses. It interprets one missing coordinate; it does not by itself prove existence, stability, or branch-wide regularity.

The 3SEG reveals the dimensional boundary. On a regular one-event section, returning the two carrier velocities leaves three endpoint coordinates; one value of $E_{\perp}$ therefore defines generically a two-dimensional level surface rather than a unique endpoint. Full vector return must therefore be solved first. The in-phase and anti-phase observables classify returned finite-amplitude cycles continued from the two linear seeds; they do not replace return and do not imply exact 0 or π phase locking away from the linear limit. Integrated exchange then follows as an endpoint balance and a diagnostic, not as an additional independent selector.

8.5 From passive persistence to resonant maintenance

Positive dissipation changes the total-energy statement to

$$\dot{E} = -P_{\text{diss}} + P_{\text{in}}, \quad \int_0^T P_{\text{in}} \, dt = \int_0^T P_{\text{diss}} \, dt \quad (8.6)$$

for every exact synchronized periodic response. The equality is a necessary cycle identity; it does not select a waveform or establish stability. A sinusoidal joint torque supplies the missing work and a forcing clock, while the nonlinear constrained mechanics determines the response harmonics, carrier motion, and reconstructed pose.

The mechanics-preserving bridge makes this transition without changing the conservative model. Its generalized nonconservative term changes from $Q = 0$ to Q_{loss} and then to $Q_{\text{loss}} + Q_{\text{act}}$ through

$$\mu_{\text{visc}}(\lambda) = \lambda c_*, \quad u(t; \lambda) = \lambda A_* \sin(2\pi t/T_{\lambda} + \phi_{\lambda}). \quad (8.7)$$

One synchronized branch connects the conservative anchor to the fully forced point while retaining geometry, inertia, constraint, and potential. Its return, waveform, and work checks establish a local numerical connection for one cycle, not persistence of the complete conservative family.

The broad forced-response computation answers a separate question. It keeps the two-segment kinematic architecture but changes the parameter vector and elastic law. In the orbit–forcing space (z_0, A, T) , regular zeros of the four-dimensional stroboscopic return residual form locally two-dimensional solution sheets wherever the corresponding Jacobian has full rank. The regular local components whose reconstructed pose increment is nonidentity form the Resonant Locomotion Manifold (RLM). Phase, peak and mean speed, harmonic content, and reconstructed pose define an observable map of those sheets. The observable image may fold or self-intersect and is not itself the definition of the RLM.

The quadrature result is physical but local. In the declared output-minus-input convention, a deformation fundamental lagging the torque by $\pi/2$ has a velocity fundamental aligned with the torque and therefore receives positive mean work. This harmonic argument explains the sign of the phase marker. Exact input–loss balance still uses the complete waveform and all dissipative channels. In the computed data, quadrature follows the positive-peak-speed ridge only on the lower-amplitude segment; at larger input it lies on a slower response or is absent from the sampled curve. It is therefore neither a universal definition of resonance nor a global optimality condition.

8.6 What local regulation adds

The forced solution sheet is a design object, not a controller. Coexisting responses can share nearly the same phase while producing different mean speeds. The control result consequently assigns one task to each observable: forcing frequency regulates the fast output–input phase error, forcing amplitude changes slowly according to mean body-frame longitudinal speed, and stage guards reject responses with unacceptable phase, deformation, yaw, or speed progress. Ordered forcing-amplitude levels are used only after a run for interpretation.

The two-loop computation moves phase and mean speed across the supported rising observable interval from two named initializations and reaches terminal mean speeds near 12.068 with phase-error root-mean-square values near 1.1×10^{-3} rad. Separate phase-only trials show why the speed coordinate is necessary: a small phase error can remain locked to an unintended low-speed response. Fourteen sampled $K_G = 5$ reduced controller matrices have negative spectral abscissa, which supports the local design on the first fourteen ordered forcing-amplitude levels. These finite linearizations do not transfer stability to the exact nonlinear plant. The nominal campaigns contain no rejection. A separate harness constructs one rejected proposal, reuses an explicit copy of the last accepted state, and accepts a reduced proposal. It tests the simulator and guard but not the rejection branch of the production stage manager, general recovery, or a global acquisition basin.

Figure 8.1 summarizes the links that are established and those that remain open without merging the two forced constructions.




Figure 8.1: Scientific topology of the thesis. Solid arrows connect results established at the analytical or numerical level stated in the text. The arrow from forced responses to control denotes use of offline observables for design and interpretation, not online membership tracking. The dashed diagonal does not identify the mechanics-preserving branch with the parameter-distinct forced-response sheets; extending that local bridge across family and parameter coordinates remains an open problem.

8.7 Limits that define the next questions

The first open question is geometric. The computed conservative and forced branches are compatible with smooth local manifold charts, but global smoothness across bifurcations, modal sectors, parameter changes, and mechanism dimensions has not been established. Scaled return Jacobians, bordered continuation singular values, section transversality, and identity tracking would turn the local manifold interpretation into a numerical regularity assessment wherever the stored data permit it.

The second question is dynamical. The mechanics-preserving bridge should be extended from one anchor across conservative family coordinates, while folds and branch loss are detected rather than assumed away. Floquet analysis, perturbed integrations, and systematic end-to-end rejection and recovery trials are needed before local regulation can become a persistence or robustness result. The relation between that bridge and the parameter-distinct response sheets, including any regular transported RLM components, must be tested after the mechanical parameters and elastic laws have been aligned.

The third question is physical. Identified damping, actuator limits, sensing noise, and hardware experiments must replace the present ideal coefficients. Hybrid contact introduces impacts and switching constraints; fluid interaction introduces environmental energy and momentum exchange. Both extensions may preserve the guiding separation between internal return and open pose, but they change the power balance and cannot be covered by the present smooth constant-rank model without new derivations (Brogliato, 2016; Prasad et al., 2023).

8.8 Conclusion

The thesis develops a route from natural oscillation to natural locomotion. Linear modes provide seeds, continuation reveals finite-amplitude families, reduced return keeps the internal dynamics repeatable while leaving pose open, and the effective oscillator–propulsive channel split gives the remaining return condition a mechanical meaning. The scalar 2SEG admits an exact exchange–return equivalence; the 3SEG shows why several transverse coordinates require vector return and modal identity.

When losses are added, accompaniment replaces passive persistence. Input work maintains synchronized nonlinear responses; where the return equations are regular, their solution sheets provide local RLM charts. Phase–speed feedback then regulates those observables across one supported interval without prescribing its complete choreography. The enduring principle is therefore simple:

let the locomotor reveal its intrinsic families, then supply and regulate only what is needed to accompany them.

Appendix A

A common reduction for continuously constrained locomotors

The two- and three-segment locomotors differ in dimension, but their equations are obtained by the same operation. This appendix presents that operation once, then identifies the model-specific ingredients used in [Chapters B and D](#). The derivation applies to smooth, bilateral, ideal velocity constraints that remain continuously active and have constant rank. It does not cover impact resets, unilateral contact, stick-slip transitions, or changes of contact mode.

A.1 Admissible velocities and virtual displacements

Let $q \in Q$ be an n -dimensional configuration and consider k independent linear velocity constraints

$$A(q)\dot{q} = 0, \quad \text{rank } A(q) = k. \quad (\text{A.1})$$

The admissible velocity space $\mathcal{D}_q = \ker A(q) \subset T_q Q$ therefore has dimension $m = n - k$. Choose a smooth full-column-rank matrix $H(q) \in \mathbb{R}^{n \times m}$ whose image is this kernel:

$$\dot{q} = H(q)\eta, \quad \text{Im } H(q) = \ker A(q), \quad A(q)H(q) = 0. \quad (\text{A.2})$$

The components of η are admissible quasi-velocities. They are not new configuration coordinates in general: for a nonintegrable distribution there need not exist a change of configuration variables whose time derivative is η .

The same distribution restricts ideal virtual displacements to $\delta q = H(q)\delta\zeta$. Thus the matrix that parameterizes admissible velocities also supplies the directions on which the dynamical balance must hold.

A.2 Projected Lagrange–d’Alembert equation

For a Lagrangian $L(q, \dot{q}) = T(q, \dot{q}) - V(q)$, applied generalized force $Q(q, \dot{q}, t)$, and constraint multiplier λ , the constrained balance is

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q = A(q)^\top \lambda. \quad (\text{A.3})$$

Multiplication by $H(q)^\top$ removes the unknown reaction because $H^\top A^\top = 0$:

$$H(q)^\top \left[\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q \right] = 0. \quad (\text{A.4})$$

After substituting $\dot{q} = H(q)\eta$, this equation has the reduced form

$$M_r(q)\dot{\eta} + h_r(q, \eta) + H(q)^\top \nabla V(q) = H(q)^\top Q, \quad M_r = H^\top M_q H. \quad (\text{A.5})$$

Here M_q is the configuration-space inertia matrix. The vector h_r collects the velocity-quadratic terms generated both by $M_q(q)$ and by the configuration dependence of $H(q)$. Computing M_r alone is consequently not enough: the projected inertial vector must be derived with the same velocity convention.

[Figure A.1](#) summarizes the resulting sequence from the mechanism description to reduced integration and pose reconstruction.




Figure A.1: Common modelling sequence for the two locomotors. The constraint matrix and kernel basis change with the mechanism, but projection eliminates the reaction multiplier in the same way. Pose reconstruction is the last operation; it is not another name for constraint reduction.

Table A.1: The same projected construction in the two locomotors.

	Two-segment locomotor	Three-segment locomotor
Configuration	$q = (x, y, \theta, \delta), n = 4$	$q = (x, y, \theta, \delta_1, \delta_2), n = 5$
Constraint rank	one no-side-slip row	one no-side-slip row
Admissible velocity	$\eta = (v, \omega, \sigma), m = 3$	$\xi = (v, \omega, \sigma_1, \sigma_2), m = 4$
Kernel basis	H in Equation (3.3)	N in Equation (D.2)
Transverse oscillatory directions	one after carrier elimination	two after carrier elimination
Return implication	a scalar section-storage value can identify the missing crossing speed	after a regular one-event section and carrier return, three endpoint coordinates remain; one storage value leaves generically a two-dimensional level surface, so full vector return remains necessary

A.3 Why ideal constraints redistribute but do not energize

Let $E = \dot{q}^T \partial_{\dot{q}} L - L$ be the mechanical energy. For a time-independent Lagrangian, multiplication of Equation (A.3) by $\dot{q}^T$ gives

$$\dot{E} = \dot{q}^T Q + \underbrace{\dot{q}^T A(q)^T \lambda}_{(A(q)\dot{q})^T \lambda = 0}. \quad (\text{A.6})$$

The ideal constraint reaction therefore has zero total power. This does not make the constraint mechanically irrelevant. By selecting $\mathcal{D}_q$ and entering the reduced kinetic metric M_r , it structures how motion in one admissible direction couples to another. The effective oscillator–propulsive–channel split in Chapter 4 is a decomposition of that conservative coupling, not an external energy source.

When losses and actuation are introduced, only their applied powers appear:

$$Q = Q_{\text{loss}} + Q_{\text{act}}, \quad \dot{E} = \dot{q}^T Q_{\text{loss}} + \dot{q}^T Q_{\text{act}}. \quad (\text{A.7})$$

This identity is the common mechanical origin of the cycle-integrated input–loss balance used in Chapter 6.

A.4 The two model instances

Table A.1 separates three dimensions that can otherwise be confused: configuration dimension, admissible-velocity dimension, and the number of transverse oscillatory directions exposed after carrier elimination. The last number is a property of the derived decomposition; it is not the total degree of freedom of the locomotor.

For the two-segment locomotor, the dimensional construction and normalization are given in Chapter B. For the three-segment locomotor, Section D.1 instantiates the same projection with N in place of H . In both cases the pose variables are absent from the reduced field after expressing translational velocities in the body frame. Once a reduced trajectory has been obtained, the group equation is integrated to recover the pose. That final integration is the pose reconstruction used throughout the manuscript.

A.5 Scope of the shared method

The projection in Equation (A.4) is general within its stated regularity class and is standard in nonholonomic mechanics (Bloch, 2015). Its use here does not establish a theory for all forms of environmental interaction. A walking model with impacts requires guards and reset maps; unilateral contact requires complementarity or contact mode logic; fluid locomotion requires explicit fluid forces or a coupled fluid–structure model. Those systems may preserve the organizing question of natural locomotion, but they do not inherit Equation (A.5) without new modelling work.

Appendix B

Two-segment model and conventions

This appendix gives the dimensional origin of the reduced equations used in [Chapters 3 and 4](#). Its purpose is traceability: the main text works with the compact dimensionless system, while the equations below show which physical balance produces each coefficient and which sign convention is being used.

B.1 Geometry, velocities, and ideal constraint

Let m_1, I_1 be the mass and planar yaw inertia of the carrier and let m_2, I_2 denote the corresponding quantities for the internal segment. The distance from the contact point to the carrier centre of mass is l_1 ; the distance from the joint to the internal-segment centre of mass is l_2 . The configuration and admissible velocity coordinates are

$$q = (x, y, \theta, \delta), \quad \eta = (v, \omega, \sigma)^\top, \quad \dot{q} = H(\theta)\eta, \quad (\text{B.1})$$

with H given in [Equation \(3.3\)](#). Thus v is positive along the carrier heading, $\omega = \dot{\theta}$, and $\sigma = \dot{\delta}$. The absolute angular speed of the internal segment is $\omega + \sigma$. This last convention fixes the signs of all mixed inertia terms.

Let $r_C = (x, y)^\top$ be the constrained contact point and define

$$e(\varphi) = \begin{bmatrix} \cos \varphi \\ \sin \varphi \end{bmatrix}, \quad e_\perp(\varphi) = \begin{bmatrix} -\sin \varphi \\ \cos \varphi \end{bmatrix}. \quad (\text{B.2})$$

In the geometry used by the reduced model, the carrier centre of mass and the joint are co-located at

$$r_1 = r_J = r_C + l_1 e(\theta), \quad r_2 = r_J + l_2 e(\theta + \delta), \quad (\text{B.3})$$

where r_2 is the internal-segment centre of mass. Since $\dot{r}_C = v e(\theta)$,

$$\dot{r}_1 = v e(\theta) + l_1 \omega e_\perp(\theta), \quad (\text{B.4})$$

$$\dot{r}_2 = \dot{r}_1 + l_2 (\omega + \sigma) e_\perp(\theta + \delta). \quad (\text{B.5})$$

These positions remove any ambiguity about the reference points of l_1 and l_2 .

The constrained Lagrange-d'Alembert equation is

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = A(\theta)^\top \lambda, \quad L = T - U(\delta). \quad (\text{B.6})$$

Because $A(\theta)H(\theta) = 0$, multiplication by $H^\top$ eliminates the reaction multiplier. The ideal lateral reaction consequently performs no work on an admissible velocity. The projected balance is

$$M_{\text{dim}}(\delta) \dot{\eta} + h_{\text{dim}}(\delta, \eta) + \begin{bmatrix} 0 & 0 & U'(\delta) \end{bmatrix}^\top = 0. \quad (\text{B.7})$$

B.2 Dimensional reduced equations

Set $M_{\text{tot}} = m_1 + m_2$ and $D = \omega + \sigma$. The kinetic energy is

$$T = \frac{1}{2} m_1 \|\dot{r}_1\|^2 + \frac{1}{2} I_1 \omega^2 + \frac{1}{2} m_2 \|\dot{r}_2\|^2 + \frac{1}{2} I_2 D^2. \quad (\text{B.8})$$

Direct evaluation and projection give

$$M_{\text{dim}}(\delta) = \begin{bmatrix} M_{\text{tot}} & -m_2 l_2 \sin \delta & -m_2 l_2 \sin \delta \\ -m_2 l_2 \sin \delta & I_1 + I_2 + M_{\text{tot}} l_1^2 + m_2 l_2^2 + 2m_2 l_1 l_2 \cos \delta & I_2 + m_2 l_2^2 + m_2 l_1 l_2 \cos \delta \\ -m_2 l_2 \sin \delta & I_2 + m_2 l_2^2 + m_2 l_1 l_2 \cos \delta & I_2 + m_2 l_2^2 \end{bmatrix}. \quad (\text{B.9})$$

The quadratic inertial vector is

$$h_{\text{dim}} = \begin{bmatrix} -M_{\text{tot}}l_1\omega^2 - m_2l_2 \cos \delta D^2 \\ (M_{\text{tot}}l_1 + m_2l_2 \cos \delta)v\omega - m_2l_1l_2 \sin \delta (D^2 - \omega^2) \\ m_2l_2 \cos \delta v\omega + m_2l_1l_2 \sin \delta \omega^2 \end{bmatrix}. \quad (\text{B.10})$$

The first row is the longitudinal balance, the second is the carrier-yaw balance, and the third is the internal-joint balance. In particular, the finite-speed product $v\omega$ is present in the mechanical model before the change of variables $q_y = v\omega$ used in Chapter 4.

The signs in Equations (B.9) and (B.10) can be checked without redoing the full variational calculation. At $\delta = 0$ the translation–rotation inertia terms proportional to $\sin \delta$ vanish. At $\omega = \sigma = 0$ the entire drift vector vanishes. Hence every state $(0, v, 0, 0)$ is a reduced equilibrium, exactly as required by Equation (3.17).

B.3 Dimensionless dictionary and energy check

Let

$$L = l_1 + l_2, \quad \alpha = \frac{l_1}{L}, \quad \beta = \frac{l_2}{L}, \quad \alpha + \beta = 1, \quad (\text{B.11})$$

and let t_* be the reference time. Hats denote dimensionless variables:

$$\hat{x} = \frac{x}{L}, \quad \hat{t} = \frac{t}{t_*}, \quad \hat{v} = \frac{t_*v}{L}, \quad \hat{\omega} = t_*\omega, \quad \hat{\sigma} = t_*\sigma, \quad \hat{E} = \frac{t_*^2 E}{M_{\text{tot}}L^2}. \quad (\text{B.12})$$

Because the joint angle δ is already dimensionless, dimensional spring coefficients in $U_{\text{dim}} = \frac{1}{2}k_{2,\text{dim}}\delta^2 + \frac{1}{4}k_{4,\text{dim}}\delta^4$ become

$$\hat{k}_2 = \frac{t_*^2 k_{2,\text{dim}}}{M_{\text{tot}}L^2}, \quad \hat{k}_4 = \frac{t_*^2 k_{4,\text{dim}}}{M_{\text{tot}}L^2}. \quad (\text{B.13})$$

The numerical chapters drop hats after this normalization. Their periods, velocities, and energies are therefore dimensionless unless an explicit dimensional calibration is supplied; in particular, the reported period range and the values $k_2 = 1$, $k_4 = 10$ belong to this numerical scale. When $k_{2,\text{dim}} > 0$, choosing $t_* = \sqrt{M_{\text{tot}}L^2/k_{2,\text{dim}}}$ sets $\hat{k}_2 = 1$; the computations may equivalently be read as fixing that reference clock.

The remaining design ratios are

$$\mu = \frac{m_2}{M_{\text{tot}}}, \quad j_i = \frac{I_i}{M_{\text{tot}}L^2}, \quad (\text{B.14})$$

and the design used in the conservative computations specializes to

$$\alpha = \varepsilon, \quad \beta = 1 - \varepsilon, \quad \mu = 1 - \varepsilon, \quad j_1 = \gamma\alpha^3, \quad j_2 = \gamma\beta^3. \quad (\text{B.15})$$

Dividing Equation (B.9) by the common mass–length scales then produces

$$\frac{m_2l_2}{M_{\text{tot}}L} = \mu\beta, \quad \frac{I_1}{M_{\text{tot}}L^2} = j_1, \quad \frac{I_2}{M_{\text{tot}}L^2} = j_2, \quad (\text{B.16})$$

and hence the functions $\rho, B_{11}, B_{12}, B_{22}$ in Equation (3.12). The common minus sign in the first row and column of Equation (3.11) is inherited from the chosen positive direction of δ ; it is not absorbed into the definition of ρ .

For the unforced conservative model, the projected energy is

$$E = \frac{1}{2}\eta^\top M(\delta)\eta + U(\delta). \quad (\text{B.17})$$

Differentiating and substituting Equation (3.10) gives

$$\dot{E} = \eta^\top \left(M\dot{\eta} + h + \begin{bmatrix} 0 & 0 & U' \end{bmatrix}^\top \right) = 0, \quad (\text{B.18})$$

where the quadratic vector is the one induced by the configuration-dependent kinetic metric. The identity is analytical; the population audit below reports only the energy variation resolved by the 40 stored phase samples and does not turn that diagnostic into a raw collocation residual. It concerns total mechanical energy. The Schur-complement storage $E_\perp$ in Chapter 4 is a different scalar: its derivative need not vanish after the propulsive channel is reopened.

B.4 Population-level numerical checks

The displayed continuation window contains 379 exported cycles with $|\bar{v}| \leq 200$. Applying the same definitions to every stored member gives the summary in Table B.1.

The zero endpoint defect states that the first and last exported samples coincide; it is not a substitute for the unavailable raw collocation residual or an independent exact-field integration. The small continued value of κ is consistent with

Table B.1: Checks over the complete 379-cycle two-segment window. Endpoint and energy quantities are evaluated on the 40 stored phase samples of each cycle; pose is reconstructed by composite trapezoidal integration.

Quantity	Value over the stored population
Mean-speed interval	$[-199.260, 199.067]$
Period interval	$[0.66740, 0.79877]$
Maximum stored reduced-state endpoint defect	0
Maximum relative stored energy span	5.19×10^{-4}
Maximum absolute continued κ	7.79×10^{-12}
Minimum reconstructed translation norm	8.86×10^{-2}

the exact identity in Equation (3.26). All 379 reconstructed pose increments are numerically nonidentity in this window. The smallest translation occurs near $\bar{v} = 0.133$ and remains 8.86×10^{-2} in the chosen length scale. This finite sample establishes the stated stored observables and nonidentity pose increments, but not population-wide dynamical qualification, uncomputed points, or regularity between samples. The corresponding post-processing is reproducible from the promoted cycle surface; independent exact-field and rank audits apply to the 29 matched roots of Chapter C.

B.5 Conventions carried into the numerical chapters

The calculations use the following conventions throughout:

1. exact reduced return is imposed on $(\delta, v, \omega, \sigma)$, while (x, y, θ) is reconstructed afterward;
2. a full-cycle section at $\delta = 0$ includes a prescribed sign of σ ;
3. $\bar{v} = T^{-1} \int_0^T v \, dt$ is a branch observable, not a replacement for pose reconstruction;
4. the scalar theorem uses $q_y = v\omega$ only on a chart where the forward speed stays away from zero;
5. damping and joint torque are absent in Chapters 3 and 4; when introduced later, their power terms appear explicitly in the energy balance.

Appendix C

Scalar field, proof details, and numerical protocol

This appendix records the finite-speed equations and numerical acceptance checks used for the 29 exchange–return rows of [Chapter 4](#). The metric derivative is retained in the exchange power and is not inserted into the integrated opened shape field.

C.1 Coefficient functions and opened field

For the dimensionless parameters introduced in [Chapters 3 and 4](#), define

$$c_{\text{geom}}(\delta) = \alpha + \mu\beta \cos \delta, \quad (\text{C.1})$$

$$A_q(\delta) = \alpha j_2 + \alpha\beta^2 \mu(1 - \mu) - \mu\beta j_1 \cos \delta, \quad (\text{C.2})$$

$$r(\delta) = \frac{A_q(\delta)}{\Delta(\delta)}, \quad (\text{C.3})$$

$$E_{\text{geom}}(\delta) = B_{11}(\delta)(c_{\text{geom}}(\delta) - \alpha) + \alpha\rho(\delta)^2. \quad (\text{C.4})$$

With the scaled state

$$y = (\delta, \sigma, \nu, q_y), \quad h = |\bar{v}|^{-1}, \quad \nu = hv, \quad q_y = v\omega, \quad (\text{C.5})$$

the physical velocities are

$$v = \frac{\nu}{h}, \quad \omega = \frac{hq_y}{\nu}. \quad (\text{C.6})$$

The chart requires $h > 0$ and $\nu \neq 0$. The shape acceleration is

$$F_\sigma = \frac{r\{q_y - \rho(\sigma + \omega)^2\} - U'}{M_\perp}. \quad (\text{C.7})$$

The two transport rows are evaluated in a fixed dependency order. First set

$$\lambda_{\text{tr}} = \frac{c_{\text{geom}}}{\Delta} (\nu - 2h\rho\sigma), \quad (\text{C.8})$$

$$\phi_{\text{tr}} = \nu \left[-\frac{B_{12} - \rho^2}{\Delta} F_\sigma + \frac{\rho c_{\text{geom}}}{\Delta} \sigma^2 \right], \quad (\text{C.9})$$

$$R_{\text{tr}} = \frac{\rho(B_{11} - B_{12})F_\sigma + E_{\text{geom}}(2\omega\sigma + \sigma^2) + B_{11}c_{\text{geom}}\omega^2}{\Delta}. \quad (\text{C.10})$$

Then

$$F_q = \frac{\phi_{\text{tr}} + h^2\nu^{-1}q_y R_{\text{tr}} - \lambda_{\text{tr}}q_y}{h}, \quad (\text{C.11})$$

and

$$F_\nu = \frac{h\rho F_\sigma + h(c_{\text{geom}} - \alpha)\sigma^2 + h^2\rho\nu^{-1}F_q + 2h^2(c_{\text{geom}} - \alpha)\nu^{-1}q_y\sigma + h^3c_{\text{geom}}\nu^{-2}q_y^2}{1 + h^2\rho\nu^{-2}q_y}. \quad (\text{C.12})$$

The denominator

$$D_{\text{fs}} = 1 + h^2\rho\nu^{-2}q_y \quad (\text{C.13})$$

is monitored along every returned orbit.

The exact scaled system is

$$\dot{y} = (\sigma, F_\sigma, F_\nu, F_q). \quad (\text{C.14})$$

For $s = \text{sgn}(\bar{v})$, the oriented variables

$$Y = (\delta, \zeta, u, Q), \quad (\sigma, \nu, q_y) = (s\zeta, su, sQ) \quad (\text{C.15})$$

and time change $d/d\vartheta = s d/dt$ give

$$\frac{dY}{d\vartheta} = (\zeta, F_\sigma, F_\nu, F_q), \quad (\text{C.16})$$

where the three field components on the right are evaluated at the signed scaled state. This is an exact conjugacy, not a large-speed approximation.

C.2 Exchange identity in the implemented convention

The conservative transverse field is

$$F_{\text{cl}} = -\frac{U' + \frac{1}{2}M'_\perp\sigma^2}{M_\perp}, \quad (\text{C.17})$$

whereas the opened field is Equation (C.7). Their difference is

$$M_\perp(F_\sigma - F_{\text{cl}}) = r\{q_y - \rho(\sigma + \omega)^2\} + \frac{1}{2}M'_\perp\sigma^2. \quad (\text{C.18})$$

Consequently,

$$P_{\text{ex}} = \sigma M_\perp(F_\sigma - F_{\text{cl}}) = \sigma \left[r\{q_y - \rho(\sigma + \omega)^2\} + \frac{1}{2}M'_\perp\sigma^2 \right]. \quad (\text{C.19})$$

Direct differentiation of

$$E_\perp = \frac{1}{2}M_\perp(\delta)\sigma^2 + U(\delta) \quad (\text{C.20})$$

along the opened field gives

$$\frac{dE_\perp}{dt} = \sigma \left(M_\perp F_\sigma + U' + \frac{1}{2}M'_\perp\sigma^2 \right) = P_{\text{ex}}. \quad (\text{C.21})$$

Under the oriented time transformation, the accumulator uses

$$\frac{dI_{\text{ex}}}{d\vartheta} = sP_{\text{ex}}, \quad (\text{C.22})$$

so it integrates exchange along the physical-time traversal represented by the oriented orbit. For $s = +1$ this is the forward-period integral. For $s = -1$, increasing ϑ traverses physical time backward, and the result is the backward-period integral. Reversing a complete periodic traversal changes its sign but not its zero set; this is the only orientation-invariant fact used by the return theorem.

C.3 Event integration and nonlinear solve

For each requested $\bar{v}$, a regular closed-support cycle supplies a period scale and a section seed. The numerical unknowns are

$$p = (a, b, Q_0), \quad Y_0(p) = (0, e^a, e^b, Q_0). \quad (\text{C.23})$$

The integration state appends the exchange and speed accumulators,

$$\hat{Y} = (Y, I_{\text{ex}}, I_u), \quad \frac{dI_u}{d\vartheta} = u. \quad (\text{C.24})$$

After an initial dead interval equal to 5% of the support-period scale, event detection terminates at the next crossing $\delta = 0$ with positive direction. The reproducible continuation solver defaults to the eighth-order explicit Runge–Kutta integrator DOP853 (Hairer et al., 1993, Section II.5), relative tolerance 2×10^{-9} , absolute tolerance 2×10^{-11} , and a maximum step equal to one support period divided by 500. The integration window extends to eight support periods. The exact option set used in the original 29-row computation was not retained. The validation below therefore rests on independent re-integration of the stored initial states, rather than on an assumed historical tolerance setting.

At return, the residual is

$$\mathcal{R}_{\text{ex}} = \begin{bmatrix} I_{\text{ex}}(T)/e_{\text{sc}} \\ \{u(T) - u_0\}/u_{\text{sc}} \\ \{Q(T) - Q_0\}/Q_{\text{sc}} \\ I_u(T)/T - 1 \end{bmatrix}, \quad (\text{C.25})$$

with

$$e_{\text{sc}} = \max(1, |I_{\text{ex}}|, |\Delta E_\perp|), \quad u_{\text{sc}} = \max(1, |u_0|), \quad Q_{\text{sc}} = \max(1, |Q_0|). \quad (\text{C.26})$$

Table C.1: Independent re-integration and regular-domain checks for the 29 exchange–return rows. State defects compare the re-integrated terminal state with the stored initial state.

Quantity	Observed bound or range
successful intended returns	29/29
$\max \delta(T) $	4.84×10^{-15}
terminal ζ	[7.31799, 7.99433]
u over all trajectories	[0.562472, 1.438482]
D_{fs} over all trajectories	[0.963830, 1.101637]
$\max \delta $ over all trajectories	0.993057
$\min \Delta$ over all trajectories	1.43675
$\min M_{\perp}$ over all trajectories	0.0914364
maximum terminal-state defect	9.86×10^{-12}

The four rows are retained for three free initial coordinates, but conservation prevents them from being four independent conditions. At a returned point the first three rows describe a fixed point of the conservative section map, so the total-energy first integral makes their differential rank at most two. The rank audit below resolves rank two and shows that the mean-speed row, which fixes the requested gauge, supplies the third independent direction. The continuation driver permits at most 120 function evaluations and four initial starts. The nominal acceptance threshold is

$$\|\mathcal{R}_{\text{ex}}\|_2 \leq 10^{-8}. \quad (\text{C.27})$$

The reported rows are substantially below that threshold.

C.4 Regular domain and numerical validation

A small algebraic residual is accepted only together with the following checks:

1. the next oriented event exists and has $\zeta > 0$;
2. u stays away from zero and D_{fs} stays away from zero;
3. $\Delta(\delta) > 0$ and $M_{\perp}(\delta) > 0$ along the orbit;
4. the integration remains finite and inside the declared regular corridor;
5. the returned orbit has the requested mean-speed gauge;
6. comparison with the independent continuation uses phase-invariant closed-loop geometry rather than pointwise time alignment.

The field implementation was checked against the raw four-dimensional model on nine states. The largest right-hand-side mismatch was

$$5.6843 \times 10^{-14}, \quad (\text{C.28})$$

and the largest raw-scaled coordinate round-trip error was

$$5.5511 \times 10^{-17}. \quad (\text{C.29})$$

An independent directional-derivative check of the exchange identity gave

$$\max \left| \frac{dE_{\perp}}{dt} - P_{\text{ex}} \right| = 6.94 \times 10^{-10}. \quad (\text{C.30})$$

The stored initial states were also re-integrated independently with a stricter DOP853 configuration. All 29 integrations reached the intended section. The finite margins in Table C.1 verify the chart and positivity assumptions on this computed corridor; they do not prove them on a larger domain.

The public family contains 29 comparison rows. Its aggregate metrics are

Quantity	Maximum over the 29 rows
scaled exchange-return residual	5.5074×10^{-12}
absolute reference-speed mismatch	7.3499×10^{-11}
absolute period error	1.4984×10^{-5}
absolute maximum-shape error	1.8554×10^{-5}
absolute maximum-yaw-rate error	8.3610×10^{-4}
phase-aligned scaled loop RMS	3.0698×10^{-3}

These numbers establish the local numerical realization used in [Chapter 4](#). The accepted corridor is not an exhaustive annulus, a classification of physical designs, or a stability certificate.

Scaled compatibility and local rank. Use $Q_{\text{sc}} = \max(1, |Q_0|)$, $u_{\text{sc}} = \max(1, u_0)$ and $v_{\text{sc}} = \max(1, |\bar{v}|)$. With the scales above, centred differences of step 10^{-5} , DOP853 tolerances $(10^{-11}, 10^{-13})$, maximum step $T/500$ and relative singular-value-decomposition (SVD) threshold 10^{-7} , all 29 roots are recovered with maximum residual 5.5102×10^{-12} . The return-only 3×3 block has rank two (largest null singular value 2.14×10^{-10}). The mean-speed row raises the 4×3 rank to three, with $\sigma_{\min} \in [0.021917, 0.104764]$ and $\kappa_2 \leq 48.7703$. Adding the scaled $\bar{v}$ column gives a rank-three 4×4 Jacobian: its smallest nonzero singular value is in $[0.025463, 0.107124]$, its null singular value is at most 1.64×10^{-10} , and their ratio at most 5.28×10^{-9} . Exact energy dependence and these three nonzero directions give a local one-dimensional phase-fixed branch at each sampled root, but not a certificate between samples, through $\bar{v} = 0$, of global embedding, or of stability.

Appendix D

Three-segment model and returned-cycle population

This appendix gives the exact model and the population-level verification supporting [Chapter 5](#). Every cycle file referenced by the in-phase and anti-phase branch summaries is evaluated, so the reported properties concern the complete represented population rather than selected plots. The earlier closed-support projection is not stored in those files and is therefore not reproduced population-wide.

D.1 Exact mechanical model

Let

$$e(\varphi) = \begin{bmatrix} \cos \varphi \\ \sin \varphi \end{bmatrix}, \quad e^\perp(\varphi) = \begin{bmatrix} -\sin \varphi \\ \cos \varphi \end{bmatrix}. \quad (\text{D.1})$$

The configuration is $q = (x, y, \theta, \delta_1, \delta_2)$, where $r_C = (x, y)^\top$ is the no-side-slip contact point. The constraint and an admissible velocity basis are

$$\begin{bmatrix} -\sin \theta & \cos \theta & 0 & 0 & 0 \end{bmatrix} \dot{q} = 0, \quad \dot{q} = N(q)\xi, \quad N(q) = \begin{bmatrix} \cos \theta & 0 & 0 & 0 \\ \sin \theta & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (\text{D.2})$$

with $\xi = (v, \omega, \sigma_1, \sigma_2)^\top$. Thus $\dot{x} = v \cos \theta$, $\dot{y} = v \sin \theta$, $\dot{\theta} = \omega$, and $\dot{\delta}_i = \sigma_i$.

Set $\ell_2 = L_2/2$ and $\ell_3 = L_3/2$. The three body centers are

$$\begin{aligned} r_1 &= r_C + l_1 e(\theta), \\ r_2 &= r_1 + \ell_2 e(\theta + \delta_1), \\ r_3 &= r_1 + L_2 e(\theta + \delta_1) + \ell_3 e(\theta + \delta_1 + \delta_2), \end{aligned} \quad (\text{D.3})$$

and their velocities are

$$\begin{aligned} \dot{r}_1 &= v e(\theta) + l_1 \omega e^\perp(\theta), \\ \dot{r}_2 &= \dot{r}_1 + \ell_2 (\omega + \sigma_1) e^\perp(\theta + \delta_1), \\ \dot{r}_3 &= \dot{r}_1 + L_2 (\omega + \sigma_1) e^\perp(\theta + \delta_1) \\ &\quad + \ell_3 (\omega + \sigma_1 + \sigma_2) e^\perp(\theta + \delta_1 + \delta_2). \end{aligned} \quad (\text{D.4})$$

The corresponding angular rates are $\Omega_1 = \omega$, $\Omega_2 = \omega + \sigma_1$, and $\Omega_3 = \omega + \sigma_1 + \sigma_2$. The kinetic energy used in the calculation is exactly

$$\mathcal{T} = \frac{1}{2} \sum_{i=1}^3 m_i \|\dot{r}_i\|^2 + \frac{1}{2} \sum_{i=1}^3 I_i \Omega_i^2 = \frac{1}{2} \xi^\top M(\delta) \xi. \quad (\text{D.5})$$

Together with

$$\begin{aligned} U(\delta) &= \frac{1}{2} k_{2,1} \delta_1^2 + \frac{1}{4} k_{4,1} \delta_1^4 + \frac{1}{2} k_{2,2} \delta_2^2 + \frac{1}{4} k_{4,2} \delta_2^4 \\ &\quad + \frac{1}{2} k_{12} (\delta_1 - \delta_2)^2, \end{aligned} \quad (\text{D.6})$$

Table D.1: Dimensionless numerical model-parameter vector p_* shared by the AP and IP branch chains.

Parameter	Value	Parameter	Value
m_1	6.0000000000	I_1	0.4186424184
m_2	0.3945707003	I_2	0.0060806994
m_3	0.2765544867	I_3	0.0003142956
l_1	0.2800000000	L_2	0.1500760073
L_3	0.1051882290	k_{12}	0.0011873112
$k_{2,1}$	0.6524353546	$k_{2,2}$	0.3008047507
$k_{4,1}$	5.2584240084	$k_{4,2}$	0.4491426471
u_1	0	u_2	0

this defines the Lagrangian $L = \mathcal{T} - U$ without an implicit geometric convention. The projected Lagrange–d’Alembert equations are

$$N(q)^\top \left[\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} - Q \right] = 0, \quad Q = (0, 0, 0, u_1, u_2)^\top. \quad (\text{D.7})$$

Equivalently,

$$M(\delta)\dot{\xi} + c(\delta, \xi) + \begin{bmatrix} 0 \\ 0 \\ \partial_{\delta_1} U \\ \partial_{\delta_2} U \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ u_1 \\ u_2 \end{bmatrix}. \quad (\text{D.8})$$

For completeness, the velocity-dependent vector in Equation (D.8) can be written compactly by defining

$$m_\Sigma = m_1 + m_2 + m_3, \quad a = \ell_2 m_2 + L_2 m_3, \quad D_1 = \omega + \sigma_1, \quad D_2 = \omega + \sigma_1 + \sigma_2, \quad (\text{D.9})$$

and $c_i = \cos \delta_i$, $s_i = \sin \delta_i$, $c_{12} = \cos(\delta_1 + \delta_2)$, and $s_{12} = \sin(\delta_1 + \delta_2)$. Its four components are

$$\begin{aligned} c_v &= -l_1 m_\Sigma \omega^2 - a c_1 D_1^2 - \ell_3 m_3 c_{12} D_2^2, \\ c_\omega &= \omega v(l_1 m_\Sigma + a c_1 + \ell_3 m_3 c_{12}) - l_1 a(2\omega \sigma_1 + \sigma_1^2) s_1 \\ &\quad - L_2 \ell_3 m_3 (2\omega \sigma_2 + 2\sigma_1 \sigma_2 + \sigma_2^2) s_2 \\ &\quad - l_1 \ell_3 m_3 [2\omega(\sigma_1 + \sigma_2) + (\sigma_1 + \sigma_2)^2] s_{12}, \\ c_{\sigma_1} &= \omega v(a c_1 + \ell_3 m_3 c_{12}) + \omega^2(l_1 a s_1 + l_1 \ell_3 m_3 s_{12}) \\ &\quad - L_2 \ell_3 m_3 (2\omega \sigma_2 + 2\sigma_1 \sigma_2 + \sigma_2^2) s_2, \\ c_{\sigma_2} &= \ell_3 m_3 (L_2 D_1^2 s_2 + l_1 \omega^2 s_{12} + \omega v c_{12}). \end{aligned} \quad (\text{D.10})$$

In the state order used by the cycle files,

$$\dot{z} = \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ M(\delta)^{-1} \left(\begin{bmatrix} 0 \\ 0 \\ u_1 \\ u_2 \end{bmatrix} - c(\delta, \xi) - \begin{bmatrix} 0 \\ 0 \\ \partial_{\delta_1} U \\ \partial_{\delta_2} U \end{bmatrix} \right) \end{bmatrix}. \quad (\text{D.11})$$

For the AP/IP population, $u_1 = u_2 = 0$, and the energy in Equations (D.5) and (D.6) is conserved by the exact field. The closed-support calculation evaluates the Schur complement and its shape derivatives from this same matrix M , then integrates the Euler–Lagrange field Equations (5.17) and (5.18); it is not inferred from energy conservation alone.

D.2 One numerical model-parameter vector

All AP and IP cycle files use the same model-parameter vector p_* . Its dimensionless values are listed in Table D.1; no conversion to hardware units is assigned in this calculation.

These values define the single model used for the coexistence result in Chapter 5.

The vector was selected by a finite, declared design-recipe search rather than by fitting hardware or optimizing the final plotted observables. Eight from-origin recipes varied the proximal inertia, local stiffness and coupling, masses, inertias, and link lengths. Promotion required reconstruction from a linear modal seed, return and modal-identity gates, an exchange-compatible lift, and a native continuation branchability test. The best-scoring in-phase recipe supplied the frozen mechanism; recovering the anti-phase sheet for that same mechanism required the recorded three-cover/source-sheet

chart. Both sectors were then regenerated and checked with every physical coefficient fixed to the table above. This provenance supports the same-mechanism coexistence claim, but it is a finite design search, not an exhaustive classification or a claim that p_* is optimal.

The continuation equations admitted an auxiliary isotropic damping coordinate μ_{aux} , which is not part of p_* . In the continuation field it enters exactly as

$$M(\delta)\dot{\xi} + c(\delta, \xi) + \begin{bmatrix} 0 \\ 0 \\ \partial_{\delta_1} U \\ \partial_{\delta_2} U \end{bmatrix} + \mu_{\text{aux}}\xi = 0, \quad u_1 = u_2 = 0. \quad (\text{D.12})$$

Consequently,

$$\dot{E} = -\mu_{\text{aux}} \xi^T \xi, \quad 0 = E(T) - E(0) = -\mu_{\text{aux}} \int_0^T \xi^T \xi \, dt. \quad (\text{D.13})$$

Every nonstationary periodic root therefore lies exactly on $\mu_{\text{aux}} = 0$. The largest absolute values stored by the finite continuation discretization are 2.68×10^{-12} on AP and 1.24×10^{-10} on IP. Reported interior field defects are evaluated against the conservative field, consistently with this exact identity.

D.3 Modal coordinates and numerical construction

At p_* , the carrier-closed linear matrices are

$$M_0 = \begin{bmatrix} 0.01806210949 & 0.002942788035 \\ 0.002942788035 & 0.001025996728 \end{bmatrix}, \quad (\text{D.14})$$

$$K_0 = \begin{bmatrix} 0.6536226658 & -0.001187311181 \\ -0.001187311181 & 0.3019920619 \end{bmatrix}.$$

The mass-normalized modal vectors and frequencies are

$$u_{\text{IP}} = [6.99370047 \quad 2.65567809]^T, \quad \Omega_{\text{IP}} = 5.835719714, \quad (\text{D.15})$$

$$u_{\text{AP}} = [7.41765930 \quad -42.69238689]^T, \quad \Omega_{\text{AP}} = 24.230934379. \quad (\text{D.16})$$

Their signs define the local same-direction and opposite-direction choreographies. The large Euclidean entries result from normalization in the small effective-inertia metric.

Closed-support continuation starts from one of these two vectors and solves four-component transverse return together with a phase row and an amplitude coordinate. The support is lifted with the energy-minimizing carrier velocity in Equation (5.14). Multiple shooting limits accumulated endpoint error; Hermite–Simpson collocation supplies a local defect representation when the AP curve becomes strongly deformed. Pseudo-arclength continuation then follows the returned opened curves at fixed p_* , including through folds of a displayed coordinate.

The original construction assessed field consistency by fitting a periodic cubic spline to the exported pose-reduced state, differentiating after trimming five percent of the samples at each end, and comparing with the exact field in Equation (D.11). The seam-inclusive check does not differentiate the stored trace. On every interval $[t_k, t_{k+1}]$, two fourth-order Runge–Kutta half steps of the exact field produce $\hat{z}_{k+1}$ from z_k . With

$$s_i = \max\{1, \text{ptp}(z_i)\}, \quad e_{k,i} = \hat{z}_{k+1,i} - z_{k+1,i}, \quad (\text{D.17})$$

the raw and scaled defects are

$$\epsilon_{\text{flow}} = \max_{k,i} |e_{k,i}|, \quad \epsilon_{\text{flow,sc}} = \max_{k,i} \frac{|e_{k,i}|}{s_i}. \quad (\text{D.18})$$

These are mesh-dependent local multiple-shooting defects, not a mesh-independent field norm. Their interpretation therefore retains the number and duration of stored intervals together with the refinement checks below. The last stored interval ends at the duplicated initial state and is included. The difference between one Runge–Kutta step and two half steps, divided by 15, estimates the local integration error. The worst raw-defect interval in each sector is also integrated with DOP853 (Hairer et al., 1993, Section II.5) at relative tolerance 10^{-12} , absolute tolerance 10^{-14} , and a maximum internal step of one sixteenth of the stored interval.

D.4 Cycle qualification and modal identity

The represented population was generated with two acceptance conditions. First, its periodic-spline raw-field defect had to satisfy

$$\epsilon_{\text{rhs}} < 5 \times 10^{-3}. \quad (\text{D.19})$$

The seam-inclusive local-flow check above independently qualifies consistency with the declared exact field. The historical spline threshold in Equation (D.19) remains the stated generation criterion. Second, the sampled cycle had to remain in its declared modal sector. For $j \in \{\text{IP}, \text{AP}\}$, define

$$Q_j(t) = u_j^\top M_0 \delta(t), \quad \tilde{P}_j(t) = \frac{u_j^\top M_0 \sigma(t)}{|\Omega_j|}, \quad (\text{D.20})$$

$$A_j = \left[\frac{1}{T} \int_0^T \left(Q_j^2 + \tilde{P}_j^2 \right) dt \right]^{1/2}. \quad (\text{D.21})$$

The expected-sector share and ratio are

$$S_j = \frac{A_j}{A_{\text{IP}} + A_{\text{AP}}}, \quad R_j = \frac{A_j}{A_{j'}}, \quad j' \neq j. \quad (\text{D.22})$$

For either position or rate signals $x(t), y(t)$, the reported correlation is the time-weighted Pearson coefficient

$$\text{corr}_T(x, y) = \frac{\int_0^T (x - \bar{x})(y - \bar{y}) dt}{\sqrt{\int_0^T (x - \bar{x})^2 dt \int_0^T (y - \bar{y})^2 dt}}, \quad \bar{x} = \frac{1}{T} \int_0^T x dt. \quad (\text{D.23})$$

The in-phase acceptance test requires both correlations to be at least 0.5, $S_{\text{IP}} \geq 0.65$, and $R_{\text{IP}} \geq 2$. The anti-phase test uses correlations at most -0.5 , $S_{\text{AP}} \geq 0.65$, and $R_{\text{AP}} \geq 2$. Recalculation by trapezoidal time integration gives 2,278/2,278 AP passes and 1,358/1,358 IP passes. The ratio condition already implies a share of at least $2/3$, so the share row records the implemented criterion rather than adding an independent restriction.

These tests classify the sampled modal content; they are not a proof that an entire invariant modal manifold exists. The linear parent vectors supply the coordinates, while the activity and correlation metrics are evaluated on the finite-amplitude cycles themselves.

D.5 Population verification

Each cycle file referenced by the branch summary is read independently. For a stored state sequence z_k at strictly increasing times t_k , the verification evaluates

$$\epsilon_{\text{ret}} = \|z_N - z_0\|_\infty, \quad (\text{D.24})$$

$$\Delta\theta = \sum_{k=0}^{N-1} \frac{\omega_k + \omega_{k+1}}{2} (t_{k+1} - t_k), \quad (\text{D.25})$$

$$\Delta x = \int_0^T v(t) \cos \theta(t) dt, \quad \Delta y = \int_0^T v(t) \sin \theta(t) dt. \quad (\text{D.26})$$

The trapezoidal pose reconstruction uses $\theta(0) = 0$; a different initial pose rigidly transforms the path but does not change whether the translation is zero. File existence, sample count, time monotonicity, endpoint return, pose reconstruction, time-weighted modal identity, total energy, and the local-flow defects in Equation (D.18) are recomputed. Cycle-mean exchange and the periodic-spline field diagnostic are also retained from the branch summary that generated the population. Table D.2 reports the resulting population-wide bounds.

The zero endpoint and stored-energy entries mean equality at the precision written in the cycle files; they are not exact-arithmetic statements. The independent DOP853 checks on the worst raw interval give defects 1.435×10^{-7} for AP and 5.423×10^{-6} for IP, differing from the two-half-step endpoints by at most 2.73×10^{-10} and 1.33×10^{-7} , respectively. Thus the independently evaluated field, not merely the duplicated endpoint, supports the returned-cycle interpretation at the reported numerical resolution. The earlier periodic-spline metric remains below its original 5×10^{-3} gate.

This audit establishes local multiple-shooting consistency of the complete represented population. It does not test attraction, reachability, or whether every stored return time is the minimal period. These properties are not used in the finite coexistence statement.

D.6 What the construction tests separate

Four finite panels vary one ingredient of the construction. They do not survey all 3SEG parameter choices; they identify which shortcut failed within the tested calculations. Table D.3 compares their distinct failure signals.

The controls separate internal excursion, modal choreography, numerical chart, raw-field consistency, and reconstructed displacement. Exchange is valuable during construction even though its cycle mean is a consequence of exact full return.

Table D.2: Population verification over both continued modal populations. The local-flow audit covers every stored interval, including the closure seam. Parenthetical labels distinguish current recomputations from the earlier periodic-spline branch diagnostic.

Quantity	AP	IP
cycle files	2,278	1,358
time samples	2,189,158	1,247,630
missing files	0	0
non-monotone time grids	0	0
maximum endpoint defect (recomputed)	0	0
stored intervals checked against exact flow	2,186,880	1,246,272
maximum local-flow defect, raw (recomputed)	1.432×10^{-7}	5.557×10^{-6}
maximum local-flow defect, scaled (recomputed)	1.001×10^{-8}	1.074×10^{-6}
maximum closure-interval defect, raw (recomputed)	3.885×10^{-8}	2.832×10^{-7}
maximum relative total-energy span (recomputed)	7.131×10^{-9}	5.948×10^{-9}
time-weighted identity passes (recomputed)	2,278/2,278	1,358/1,358
minimum expected-sector share (recomputed)	0.69246	0.96335
minimum expected/opposite ratio (recomputed)	2.2516	26.2838
smallest required-sign correlation magnitude (recomputed)	0.99630	0.56425
maximum absolute cycle-mean exchange (stored)	1.692×10^{-8}	2.764×10^{-8}
maximum transverse-energy drift (stored)	0	0
maximum periodic-spline field defect (earlier diagnostic)	1.102×10^{-3}	1.502×10^{-3}
minimum reconstructed displacement magnitude	3.038×10^{-3}	5.563×10^{-2}
maximum reconstructed displacement magnitude	79.923	143.562

Table D.3: Finite controls on activity, modal identity, period representation, and exchange-assisted construction.

Control	Observed signal	Missing property
large activity	AP-oriented internal activity reached 570.31	no moving AP branch; identity or raw-field quality could fail
modal identity	7/9 sector passes in the isolated-lever panel	0/9 moving branches
cycle representation	AP support remained numerically representable	ordinary one- and three-cover charts gave 0/4 moving trials each
exchange omitted from lift solve	one IP calculation gave ten moving identity-preserving points	raw-field mismatch $\simeq 1.04 \times 10^{-1}$; no moving AP pair

D.7 Observable convention and scope

The AP continuation stores the mean-speed observable in its active continuation coordinate. The corresponding field is inactive and constant in the IP export, so IP speed statements use the mean of $v(t)$ recomputed from each cycle file. Treating the inactive IP field as a physical speed would produce discrepancies as large as 161.1 in the model coordinates.

The population therefore combines

$$\begin{aligned} &\text{same numerical model parameters} \wedge \text{declared modal identity} \wedge \text{recomputed endpoint return} \\ &\wedge \text{field consistency} \wedge \text{small exchange} \wedge \text{nonzero reconstructed displacement.} \end{aligned} \quad (\text{D.27})$$

This is the finite numerical coexistence statement used in [Chapter 5](#). The cycle-level result is independent of a one-to-one closed-support genealogy, which is unavailable for the complete set; stability is a separate question.

Appendix E

Forced-response equations and numerical verification

This appendix documents the two forced calculations in Chapter 6. The first is the local mechanics-preserving homotopy from one conservative 2SEG cycle. The second is the independent, parameter-distinct response population containing 13,782 continued cycles at 25 forcing-amplitude levels. For the latter, reported peaks are observed mesh values and quadrature locations use only piecewise-linear interpolation between adjacent continued responses.

E.1 Numerical construction of the mechanics-preserving bridge

The bridge retains the mechanical parameters

$$\varepsilon = 0.681, \quad \gamma = 2.504, \quad k_2 = 1, \quad k_4 = 10 \quad (\text{E.1})$$

of Chapter 3. Only the declared viscous and input terms are added:

$$\mu_{\text{visc}}(\lambda) = \lambda c_*, \quad u(t; \lambda) = \lambda A_* \sin\left(\frac{2\pi t}{T_\lambda} + \phi_\lambda\right), \quad c_* = 10^{-3}. \quad (\text{E.2})$$

On the numerically refined conservative anchor, the first-order forcing amplitude that can balance isotropic viscous loss is 0.03376316109. The computation uses a 10% margin,

$$A_* = 0.03713947719, \quad (\text{E.3})$$

and selects the lower of the two work-balancing phase branches. At $\lambda = 0$ the input vanishes, so its physical phase is undefined.

Time zero is fixed by $\delta(0) = 0, \sigma(0) > 0$. For every positive λ , the unknowns are $(v_0, \omega_0, \sigma_0, T_\lambda, \phi_\lambda)$; the equations are the four synchronized reduced-return rows and the fixed time-weighted mean-speed row $\bar{v} = 10.435810958679$. The drive frequency $2\pi/T_\lambda$ therefore changes with the corrected response. Nineteen prescribed homotopy values are used:

$$\lambda \in \{0, 10^{-4}, 3 \times 10^{-4}, 10^{-3}, 3 \times 10^{-3}, 5 \times 10^{-3}, 10^{-2}, 3 \times 10^{-2}, 0.06, 0.1, 0.2, \dots, 1\}. \quad (\text{E.4})$$

Every declared point was reached. The maximum strict reduced-return residual is 7.102×10^{-10} , the maximum mean-speed error is 1.075×10^{-12} , and the maximum absolute defect in the differential energy identity is 4.112×10^{-13} . The largest relative cycle work imbalance is 2.940×10^{-4} at $\lambda = 10^{-4}$, where both work terms are approximately 10^{-5} ; at $\lambda = 1$ the relative imbalance is 2.60×10^{-11} . Meshes from 501 to 4001 points and three independent integration tolerances were compared. The largest strict-ultra endpoint difference is 5.405×10^{-14} .

The normalized waveform distance compares phase-aligned $(\delta, v, \omega, \sigma)$ traces after scaling each coordinate by an anchor amplitude. It is 4.658×10^{-5} at $\lambda = 10^{-4}$ and 0.28708 at $\lambda = 1$. This establishes numerical approach to the conservative anchor along the computed branch. The same calculation provides row-level waveform, mesh, and tolerance comparisons; the result concerns this one local branch, not persistence of the whole conservative family.

E.2 Equations of the parameter-distinct forced model

Let $\eta_f = (v, \omega_b, \sigma)^\top$. The response computation uses

$$M_f(\delta) = \begin{bmatrix} 1 & -(\varepsilon_f - 1)^2 \sin \delta & -(\varepsilon_f - 1)^2 \sin \delta \\ -(\varepsilon_f - 1)^2 \sin \delta & M_{22} & M_{23} \\ -(\varepsilon_f - 1)^2 \sin \delta & M_{23} & M_{33} \end{bmatrix}, \quad (\text{E.5})$$

Table E.1: Seam-inclusive local-flow defects over the complete forced-response population. The closure column concerns the last stored interval returning to the initial state.

State coordinate	maximum raw defect	maximum scaled defect	maximum raw closure defect
δ	3.302×10^{-6}	2.680×10^{-7}	4.391×10^{-7}
v	1.660×10^{-4}	1.716×10^{-5}	6.876×10^{-5}
ω_b	1.471×10^{-4}	1.175×10^{-5}	6.679×10^{-5}
σ	4.827×10^{-4}	3.331×10^{-6}	1.305×10^{-4}
a	1.356×10^{-8}	6.781×10^{-9}	1.356×10^{-8}
b	2.888×10^{-8}	1.444×10^{-8}	2.799×10^{-8}

where

$$M_{22} = \varepsilon_f^3 \gamma_f + \varepsilon_f^3 - \gamma_f(\varepsilon_f - 1)^3 + (1 - \varepsilon_f) [-2\varepsilon_f^2 \cos \delta + 2\varepsilon_f^2 + 2\varepsilon_f \cos \delta - 2\varepsilon_f + 1], \quad (\text{E.6})$$

$$M_{23} = (\varepsilon_f - 1)^2 [-\varepsilon_f \gamma_f + \varepsilon_f \cos \delta - \varepsilon_f + \gamma_f + 1], \quad (\text{E.7})$$

$$M_{33} = -(\varepsilon_f - 1)^3 (\gamma_f + 1). \quad (\text{E.8})$$

The continued data use $\varepsilon_f = 0.7317084397592031$ and $\gamma_f = 0.736801492464307$. A numerical minimization over one full period, $\delta \in [-\pi, \pi]$, gives minimum eigenvalue 0.02561150819; the matrix is therefore uniformly positive in the tested interval to the reported numerical precision.

The code evaluates the acceleration in right-hand-side form,

$$M_f(\delta) \dot{\eta}_f = b_f(\delta, \eta_f) + B_\delta \tau_\delta - c_d \eta_f, \quad B_\delta = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^\top, \quad (\text{E.9})$$

with

$$b_f = \begin{bmatrix} \varepsilon_f \omega_b^2 + (1 - \varepsilon_f)^2 (\sigma + \omega_b)^2 \cos \delta \\ \varepsilon_f (1 - \varepsilon_f)^2 (\sigma + 2\omega_b) \sigma \sin \delta - \omega_b v [\varepsilon_f + (1 - \varepsilon_f)^2 \cos \delta] \\ -\varepsilon_f (1 - \varepsilon_f)^2 \omega_b^2 \sin \delta - (1 - \varepsilon_f)^2 \omega_b v \cos \delta - 2\delta \end{bmatrix}. \quad (\text{E.10})$$

Thus the left-hand-side convention of Equation (6.9) is exact when

$$F_f = -b_f, \quad D_f = c_d I_3. \quad (\text{E.11})$$

In particular, the third component of F_f ends in $+2\delta$. The factor two follows from the implemented potential $V_f(\delta) = \delta^2$, for which the spring torque is -2δ . Writing the vector b_f on the left with a plus sign would reverse every conservative acceleration term and is not the implemented model.

To express the sinusoidal input in autonomous form, introduce the Cartesian forcing-clock state (a, b) . Its dynamics and torque output are

$$\dot{a} = -\Omega b + a(1 - a^2 - b^2), \quad \dot{b} = \Omega a + b(1 - a^2 - b^2), \quad \tau_\delta = A_{\text{exc}} b. \quad (\text{E.12})$$

On $a^2 + b^2 = 1$, one may write $a = \cos \psi$, $b = \sin \psi$ with $\dot{\psi} = \Omega$. The last expression is therefore a unit-amplitude sinusoid scaled by A_{exc} . The implementation uses $A_{\text{exc}} b$ directly; it does not divide by $\sqrt{a^2 + b^2}$.

E.2.1 Population-wide dynamic-return audit

For all 13,782 continued cycles, the first and last recorded values agree for δ , v , ω_b , σ , a , and b at the written precision. Dynamic consistency is tested in addition to this endpoint check. On each of the 2,756,400 stored intervals, including cycle closure, two fourth-order Runge–Kutta half steps of Equations (E.9) and (E.12) predict $\hat{z}_{k+1}$ from z_k . The scaled defect divides coordinate i by $s_i = \max\{1, \text{ptp}(z_i)\}$ on that cycle. For cycle j , the continued forcing frequency used by the field is inferred from the recorded clock turn as $\Omega_j = [\text{unwrap } \psi_j(T_j) - \text{unwrap } \psi_j(0)]/T_j$, with $\psi_j = \text{atan2}(b_j, a_j)$, rather than read from a single historical archive scalar. The population maxima are listed in Table E.1. They are mesh-dependent local multiple-shooting defects, not a mesh-independent norm of field error; the interval count, Richardson estimate, and DOP853 comparison state their scale.

The one-step/two-half-step Richardson estimate is at most 9.73×10^{-7} . Tight-tolerance DOP853 integration (Hairer et al., 1993, Section II.5) of the worst raw interval gives a defect 4.837×10^{-4} and differs from the two-half-step endpoint

by 9.38×10^{-7} . The independently coded numerical field agrees with the symbolic model builder to 1.42×10^{-13} over 72 state-parameter evaluations. Finally, Simpson integration of every cycle verifies

$$\int_0^T A_{\text{exc}} b \sigma \, dt = \int_0^T c_d (v^2 + \omega_b^2 + \sigma^2) \, dt \quad (\text{E.13})$$

with maximum relative imbalance 7.583×10^{-6} and median 5.257×10^{-9} .

Unwrapping $\psi = \text{atan2}(b, a)$ gives one forcing-clock turn per mechanical return for every cycle, with maximum winding-number deviation 2.22×10^{-16} and maximum clock-radius defect 2.86×10^{-9} . Together, recorded closure and the local-flow audit support the primary 1:1 synchronized interpretation at the reported numerical resolution. They do not establish attraction, reachability, or global regularity of the solution set. The unit winding expresses synchronization, not resonance by itself: *resonant* refers to the component seeded near the small-amplitude natural period and followed through the nonlinear response ridge; response maxima and quadrature crossings are assessed separately.

All variables in this computation are nondimensional. For reference mass M_0 , length l_0 , and the coefficient k_0 in the dimensional potential $V = k_0 \delta^2$, the time scale and torque-amplitude conversion are

$$t_c = \sqrt{\frac{M_0 l_0^2}{k_0}}, \quad A_{\text{exc}} = \frac{\tau_{\text{amp}} t_c^2}{M_0 l_0^2} = \frac{\tau_{\text{amp}}}{k_0}. \quad (\text{E.14})$$

Consequently, $\tau_{\text{amp}} = k_0 A_{\text{exc}}$. The numerical color coordinate used in the response plots equals $1000 A_{\text{exc}}$; it is an equivalent N mm value only for $k_0 = 1 \text{ N m}$.

The 25 fixed nondimensional forcing-amplitude levels, reported to eight decimal places and in increasing order, are

$$\begin{aligned} \mathcal{A} = \{ & 0.00204965, 0.00329097, 0.00451922, 0.00577260, 0.00712466, \\ & 0.00865827, 0.01043898, 0.01251222, 0.01490514, 0.01762767, \\ & 0.02067394, 0.02401735, 0.02763993, 0.03148703, 0.03551893, \\ & 0.03969468, 0.04394984, 0.04821334, 0.05243581, 0.05655617, \\ & 0.06052866, 0.06426299, 0.06771534, 0.07084794, 0.07357499 \}. \end{aligned} \quad (\text{E.15})$$

An integer subscript used below only orders these physical amplitude values; it is not an additional discrete model parameter.

E.3 Phase extracted from the continued cycles

The response map stores the phase of joint deformation relative to applied joint torque. The continuation software AUTO-07P, accessed through the PyDSTool/PyCont interface (Doedel et al., 2007; Clewley, 2012), produces a mesh that is not uniform in time. Each cycle is first interpolated periodically onto 1,000 uniformly spaced times in $[0, T)$; the duplicated endpoint is excluded. For a resampled signal y , the extractor removes its mean and forms the analytic signal

$$a_y(t) = y(t) - \langle y \rangle + i \mathcal{H}[y - \langle y \rangle](t), \quad (\text{E.16})$$

where $\mathcal{H}$ denotes the discrete Hilbert transform. The cycle phase is

$$\phi_{\delta\tau} = \text{Arg} \left[\frac{1}{N} \sum_{k=1}^N \exp\{i(\arg a_\delta(t_k) - \arg a_\tau(t_k))\} \right]. \quad (\text{E.17})$$

Thus the signal order is output minus input, the wrap interval is $[-\pi, \pi)$, and quadrature is $\phi_{\delta\tau} = -\pi/2$. This is a Hilbert analytic-signal statistic. A Fourier-coefficient phase can approximate it on a nearly harmonic response, but the two definitions are not identical on a distorted waveform.

E.4 Peak-speed and quadrature extraction

For each amplitude level $A_r \in \mathcal{A}$, let

$$v_{\text{pk},rj} = \max_k v_{rj}(t_k), \quad j_{p,r} \in \arg \max_j v_{\text{pk},rj}, \quad (T_{p,r}, v_{p,r}) = (T_{rj_{p,r}}, v_{\text{pk},rj_{p,r}}) \quad (\text{E.18})$$

where k indexes the recorded within-cycle samples and j follows continuation order. Thus the stated peak is a sampled positive maximum, not a spline estimate between nodes. A quadrature crossing is accepted only when two consecutive responses bracket $-\pi/2$. If

$$d_j = \phi_{rj} + \frac{\pi}{2}, \quad d_j d_{j+1} < 0, \quad (\text{E.19})$$

Table E.2: Quadrature verification by forcing-amplitude interval.

Forcing-amplitude levels	Crossing status	Relation to the observed speed peak
$0.00204965 \leq A_{\text{exc}} \leq 0.03551893$ (15 levels)	all 15 cross $-\pi/2$	maximum linearly interpolated speed gap 0.0635625%; maximum period gap 0.0289515%
$0.03969468 \leq A_{\text{exc}} \leq 0.05655617$ (five levels)	all five cross	crossing belongs to a much slower part of the continued response, 59.19–71.84% below the observed peak
$0.06052866 \leq A_{\text{exc}} \leq 0.07357499$ (five levels)	no sampled crossing	the nearest recorded value is not promoted to a quadrature point

the interpolation fraction and crossing speed are

$$\lambda_j = \frac{|d_j|}{|d_j| + |d_{j+1}|}, \quad (\text{E.20})$$

$$T_{q,r} = (1 - \lambda_j)T_{rj} + \lambda_j T_{r,j+1}, \quad (\text{E.21})$$

$$v_{q,r} = (1 - \lambda_j)v_{\text{pk},rj} + \lambda_j v_{\text{pk},r,j+1}. \quad (\text{E.22})$$

When several crossings occur, the first in continuation order is used. No crossing is inferred when the target is absent.

The reported proximity metric is

$$\epsilon_{v,r} = 100 \frac{|v_{q,r} - v_{p,r}|}{|v_{p,r}|}, \quad (\text{E.23})$$

with the corresponding period gap defined analogously. Both numerical values are reported in percent. Table E.2 separates the three observed forcing-amplitude regimes.

The main text rounds the lowest-15-level bound upward to 0.064%. The largest phase error evaluated at an observed speed peak on those levels is 0.0452127 rad. Near-equality of peak speed and quadrature speed is therefore a local property of the rising branch, not a property of all 25 forcing-amplitude levels.

E.5 Population-wide pose reconstruction

The continued-cycle data contain the time histories needed to reconstruct pose for every response. With $\theta(0) = 0$, the calculation uses

$$\theta(t) = \int_0^t \omega_b(s) \, ds, \quad \Delta x = \int_0^T v(t) \cos \theta(t) \, dt, \quad \Delta y = \int_0^T v(t) \sin \theta(t) \, dt. \quad (\text{E.24})$$

Changing the initial pose rigidly transforms the reconstructed path and cannot change whether Δg is the identity.

Composite trapezoidal integration on the nonuniform continuation mesh gives a nonzero translation for all 13,782 responses. Repeating the reconstruction with cumulative Simpson integration for θ and Simpson integration for $(\Delta x, \Delta y)$ gives

$$3.46365 \times 10^{-7} \leq \sqrt{\Delta x^2 + \Delta y^2} \leq 9.35730, \quad \max \frac{\|\Delta p_S - \Delta p_T\|}{\|\Delta p_S\|} = 0.6776\%, \quad (\text{E.25})$$

where $\Delta p = (\Delta x, \Delta y)$ and subscripts S and T denote the Simpson and trapezoidal reconstructions. In particular, even the smallest reconstructed translation remains nonzero under both rules. The recorded population is therefore transport-bearing. Together with the population-wide local-flow audit above, this supports synchronized return and nonidentity pose at the reported numerical resolution: every recorded response lies numerically in the transported subset $\mathcal{M}_R$. The RLM name applies only to local components of this transported set that also satisfy the rank and embedding conditions stated in Section 6.4.

E.6 Nonlinear shift from the small-amplitude period

At $\delta = v = \omega_b = \sigma = 0$, the effective shape inertia after eliminating the yaw coordinate is

$$M_{\text{lin}} = M_{33}(0) - \frac{M_{23}(0)^2}{M_{22}(0)} = 0.0258224048916. \quad (\text{E.26})$$

Since the linear restoring coefficient is 2,

$$T_{\text{lin}} = 2\pi\sqrt{M_{\text{lin}}/2} = 0.7139424640185. \quad (\text{E.27})$$

The nearby value 0.7139425323 used in an earlier response calculation is the period of a finite seed cycle, not the result of this linearization. At the observed speed maximum for $A_{\text{exc}} = 0.03551893$, $T_p(0.03551893) = 0.7410295954$, giving

$$100 \frac{T_p(0.03551893) - T_{\text{lin}}}{T_{\text{lin}}} = 3.7940\%. \quad (\text{E.28})$$

At $A_{\text{exc}} = 0.07357499$ the observed peak period is 0.7411539715, a 3.8114% shift. The period shift depends only on the recorded periods and the mechanical linearization; it is independent of the Hilbert-phase and THD extraction rules. Together, these audits support the declared primary 1:1 synchronized, transport-bearing population at the reported numerical resolution; they do not establish uniqueness, attraction, global regularity, or experimental robustness.

Appendix F

Phase–speed controller and sampled finite-harmonic evidence

This appendix specifies the online estimators, the sampled phase and amplitude updates, the stage-transition tests, and the finite-harmonic matrices used in [Chapter 7](#). The rolling windows used by feedback are kept distinct from the terminal window used to accept a completed stage.

F.1 Rolling phase and speed estimates

Let $t_n = n\Delta t$, with $\Delta t = 0.006$, and let $T_n = 2\pi/|\Omega_n|$ be the current period estimate. The phase fit uses the last

$$N_{\phi,n} = \max\left\{24, \left\lfloor \frac{0.25T_0}{\Delta t} \right\rfloor, \left\lfloor \frac{T_n}{\Delta t} \right\rfloor\right\} \quad (\text{F.1})$$

samples. Over the frequencies reached here, the active term is one current forcing period. The speed estimate uses the longer window

$$N_{v,n} = \max\left\{N_{\phi,n}, \left\lfloor \frac{3T_n}{\Delta t} \right\rfloor\right\}, \quad \bar{v}_n = \frac{1}{N_{v,n}} \sum_{j=n-N_{v,n}+1}^n v_j. \quad (\text{F.2})$$

Thus phase is updated from approximately one period and mean speed from approximately three periods. Here every v_j is the longitudinal velocity component resolved in the body frame. Consequently $\bar{v}_n$ is the local speed coordinate supplied to the controller; it is not generally the norm of the net inertial translation divided by elapsed time when yaw varies.

For a signal y_i sampled at excitation phases φ_i on the phase window, define

$$X_{K_{\text{est}}} = \begin{bmatrix} 1 & \sin \varphi_1 & \cos \varphi_1 & \cdots & \sin K_{\text{est}}\varphi_1 & \cos K_{\text{est}}\varphi_1 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 1 & \sin \varphi_N & \cos \varphi_N & \cdots & \sin K_{\text{est}}\varphi_N & \cos K_{\text{est}}\varphi_N \end{bmatrix}. \quad (\text{F.3})$$

The coefficients are obtained by a linear least-squares solve,

$$\hat{\beta}_y = X_{K_{\text{est}}}^\dagger \mathbf{y} = \begin{bmatrix} \hat{c}_{y,0} & \hat{a}_{y,1} & \hat{b}_{y,1} & \cdots \end{bmatrix}^\top. \quad (\text{F.4})$$

The dagger denotes the least-squares solution; the normal-equation inverse is not formed explicitly. With $K_{\text{est}} = 3$, the fundamental phase is

$$\phi_y = \text{atan2}(\hat{b}_{y,1}, \hat{a}_{y,1}), \quad e_{\phi,n} = \text{wrap}_{[-\pi,\pi)}\left(\phi_\delta - \phi_\tau + \frac{\pi}{2}\right). \quad (\text{F.5})$$

The same excitation clock and design matrix are used for δ and τ_δ . Their order is response minus torque. Fit amplitude, residual RMS, and the condition number of $X_{K_{\text{est}}}$ diagnose a degenerate fundamental or an inadequately covered window; only the fundamental phase difference enters feedback.

F.2 Sampled phase and amplitude laws

For a scalar r , write $\text{sat}_{[\ell,u]}(r) = \min\{u, \max\{\ell, r\}\}$. Once the phase estimate is available, the integral state and the unsaturated position form of the phase PI are

$$I_{\phi,n} = I_{\phi,n-1} + \Delta t e_{\phi,n}, \quad \Omega_n^{\text{raw}} = \Omega_0 + s(K_P e_{\phi,n} + K_I I_{\phi,n}). \quad (\text{F.6})$$

Table F.1: Representative ordered forcing-amplitude levels.

Role	A_{exc}	$\bar{v}$	Ω_{exc}	$\phi + \pi/2$ [rad]
lowest stored response	0.002050	0.0215	8.8006	0.00237
low interior response	0.007125	1.7977	8.7948	-0.01262
phase-sign calibration	0.017628	6.2112	8.7439	0.000024
last linearized level	0.031487	10.8764	8.5987	0.04259
next numerical endpoint	0.035519	12.3649	8.4790	0.05127

The actual sampled update uses the equivalent velocity form and applies the frequency bound at every step:

$$\Omega_n = \text{sat}_{[5,11]}(\Omega_{n-1} + s[K_P(e_{\phi,n} - e_{\phi,n-1}) + K_I \Delta t e_{\phi,n}]). \quad (\text{F.7})$$

Away from saturation, summing Equation (F.7) gives Equation (F.6). The previous phase error is initialized to zero, and Ω_0 is the excitation frequency carried into the stage (the Ω_{free} of the condensed position form). For the output-minus-input convention in Equation (F.5), the locally verified sign is $s = +1$; the gains are $K_P = 0.04$ and $K_I = 0.08$. The saturation limits in Equation (F.7) are frequencies per model-time unit.

The amplitude update starts only after both a phase estimate and the three-period speed estimate are available. Its phase-health weight is

$$g_{\phi,n} = \begin{cases} 1, & |e_{\phi,n}| \leq 0.20, \\ (0.55 - |e_{\phi,n}|)/0.35, & 0.20 < |e_{\phi,n}| < 0.55, \\ 0, & |e_{\phi,n}| \geq 0.55. \end{cases} \quad (\text{F.8})$$

With $e_{v,n} = \bar{v}_{\text{des}} - \bar{v}_n$, the nominal commanded amplitude rate and bounded update are

$$u_{A,n} = \text{sat}_{[-r_A, r_A]}(g_{\phi,n} K_v e_{v,n} - \mathbf{1}_{\{g_{\phi,n}=0\}} \rho(A_{\text{exc},n-1} - A_{\text{min}})), \quad (\text{F.9})$$

$$A_{\text{exc},n} = \text{sat}_{[A_{\text{min}}, A_{\text{max}}]}(A_{\text{exc},n-1} + \Delta t u_{A,n}). \quad (\text{F.10})$$

The values are

$$K_v = 2.5 \times 10^{-5}, \quad r_A = 2.5 \times 10^{-4}, \quad \rho = 10^{-4}, \quad (A_{\text{min}}, A_{\text{max}}) = (10^{-4}, 8 \times 10^{-2}). \quad (\text{F.11})$$

The middle branch of g_{ϕ} progressively slows the speed adaptation; full loss of phase lock closes the speed channel and adds a weak decay toward A_{min} . As in Chapter 6, forcing amplitudes are nondimensional, and r_A is the corresponding amplitude rate per model-time unit.

F.3 Sign calibration and representative amplitude levels

The responses used for interpretation are ordered by increasing A_{exc} . Table F.1 lists the values that delimit the controller tests and the sampled finite-harmonic analysis. The ordering is not supplied to the online controller.

At the phase-sign calibration level, a -0.5% initial frequency detuning gave final-tail phase-error RMS 2.391×10^{-3} rad with $s = +1$, while $s = -1$ gave 1.2452 rad. Correct-sign starts at $+0.5\%$ and -3% detuning ended at 2.679×10^{-3} and 4.103×10^{-3} rad. By contrast, the 15-run amplitude staircase kept small phase errors but remained on the low response sheet. These two tests distinguish local phase locking from response selection.

F.4 Stage-transition logic and terminal metrics

Each proposed speed stage lasts 150 model-time units. Acceptance is evaluated on its last quarter, $\mathcal{W}_{\text{tail}} = [112.5, 150]$, rather than on either rolling feedback window. For any signal y on a window $\mathcal{W}$ of duration $|\mathcal{W}|$, define

$$\langle y \rangle_{\mathcal{W}} = \frac{1}{|\mathcal{W}|} \int_{\mathcal{W}} y(t) dt, \quad \text{RMS}_{\mathcal{W}}(y) = \sqrt{\langle y^2 \rangle_{\mathcal{W}}}, \quad \text{ptp}_{\mathcal{W}}(y) = \max_{\mathcal{W}} y - \min_{\mathcal{W}} y. \quad (\text{F.12})$$

The numerical tests use the corresponding sample mean, root mean square, and sample maximum minus minimum. Set $\bar{v}_{\text{new}} = \langle v \rangle_{\mathcal{W}_{\text{tail}}}$. A proposal passes the response-health tests when

$$\begin{aligned} \text{RMS}_{\mathcal{W}_{\text{tail}}}(e_{\phi}) &\leq 0.12, \\ \text{ptp}_{\mathcal{W}_{\text{tail}}}(\delta) &\leq 8.5, \\ |\langle \omega \rangle_{\mathcal{W}_{\text{tail}}}| &\leq 0.08, \\ \bar{v}_{\text{new}} &\geq \bar{v}_{\text{old}} - \max(0.75, 0.12|\bar{v}_{\text{old}}|). \end{aligned} \quad (\text{F.13})$$

Here ω is body yaw rate, not excitation frequency. For a requested speed $\bar{v}_{\text{des}} \geq 8$, acceptance also requires

$$|\bar{v}_{\text{new}} - \bar{v}_{\text{des}}| \leq \max(0.65, 0.14|\bar{v}_{\text{des}}|). \quad (\text{F.14})$$

The first proposed speed increment is 2, with nominal limits $0.125 \leq \Delta \bar{v}_{\text{des}} \leq 2$. After a speed-converged accepted stage it is multiplied by 1.15; after an accepted but off-command stage below the strict high-speed threshold it is multiplied by 0.70. A rejected stage restores the previous accepted plant state, excitation phase, amplitude, and frequency, and halves the increment. Once the accepted target is at least 8, the next increment is capped at 0.25. An initialization terminates after 40 proposals or when a rejected increment falls below 0.125.

Table F.2 summarizes the two realized phase-speed progressions and their terminal forcing amplitudes.

Table F.2: Guarded phase-speed progression from the two tested initializations.

Quantity	Start from rest	Start from the stored lowest-level response
proposed stages	21	21
accepted / rejected	21/0	21/0
terminal mean speed	12.0680391251	12.0677941731
terminal phase-error RMS [rad]	$1.11897244 \times 10^{-3}$	$1.10605507 \times 10^{-3}$
terminal nondimensional A_{exc}	0.0364166157	0.0364087749
A_{exc} of closest stored response	0.0355189	0.0355189
by mean speed		
online use of the response family	none	none

Across the 42 accepted stages, the maximum phase-error RMS was 2.7012×10^{-3} rad, the maximum deformation peak-to-peak was 3.5383, and the maximum absolute mean yaw rate was 4.5720×10^{-3} . Every stage therefore remained inside the declared thresholds. Because no proposal was rejected, this nominal campaign does not exercise rejection handling; the separate stress audit below constructs one rejected and one reduced proposal outside the production stage-manager branch.

F.5 Terminal-cycle, time-step, and rejection audits

The terminal speed and phase statistics do not by themselves identify a full response cycle. For a stronger a posteriori comparison, let $z_c(\varphi)$ be the final controlled cycle and $z_r^{(k)}(\varphi)$ the stored peak-response cycle of amplitude-level index k , both sampled at the same $N = 512$ forcing phases without a free cyclic shift. Here $z = (\delta, v, \omega, \sigma)$. With component scales

$$s_j = \max_{k,\ell} z_{r,j}^{(k)}(\varphi_\ell) - \min_{k,\ell} z_{r,j}^{(k)}(\varphi_\ell), \quad d_k = \left[\frac{1}{4N} \sum_{\ell=1}^N \sum_{j=1}^4 \left(\frac{z_{c,j}(\varphi_\ell) - z_{r,j}^{(k)}(\varphi_\ell)}{s_j} \right)^2 \right]^{1/2}, \quad (\text{F.15})$$

the nearest of the 25 available peak cycles is amplitude level 15. The same calculation also compares the last two controlled cycles and the two endpoints of the last cycle. These are neighborhood and near-periodicity diagnostics, not collocation residuals or membership tests for the complete synchronized-response zero set.

Table F.3 reports the nominal integration step and a replay of the final lowest-amplitude-start stage from exactly the same plant and forcing state, over the same duration, with the step halved.

Table F.3: Terminal full-cycle and time-step audit. All cycle distances are normalized according to Equation (F.15).

Quantity	$\Delta t = 0.006$	$\Delta t = 0.003$
terminal mean speed	12.0677942	12.0677663
terminal phase-error RMS [rad]	1.1060551×10^{-3}	1.1060157×10^{-3}
nearest peak-response amplitude level	15	15
nearest-cycle distance d_{15}	1.7542932×10^{-2}	1.7563618×10^{-2}
last-two-cycle RMS difference	1.1275073×10^{-4}	3.7601361×10^{-5}
one-cycle endpoint ℓ_∞ defect	2.3416024×10^{-4}	5.9930842×10^{-5}

The two complete terminal profiles differ by a normalized RMS of 1.0740512×10^{-4} . Halving the step changes mean speed by 2.7882×10^{-5} and phase-error RMS by 3.94×10^{-8} rad. The plant was integrated here by explicit fourth-order Runge-Kutta. This single-stage refinement supports the reported terminal observables; it does not cover every stage, the rest-start run, or a controller theorem.

Finally, a stress harness starts from the accepted stage at requested speed 8.2935 and mean speed 7.5801. It deliberately bypasses the nominal high-speed step cap: a direct proposal to 12 ends at mean speed 8.9980, with speed error 3.0020, and is rejected. The rejected terminal state is discarded; the harness explicitly reuses a copy of the accepted plant and

forcing state, halves the proposed increment, and accepts target 10.1468 at mean speed 8.7457, with speed error 1.4011. Thus the production simulator and guard are exercised in one constructed reject-reuse-reduced-proposal sequence, but the rejection branch of the production stage manager is not. This is neither a recovery theorem, a basin estimate, nor a test of arbitrary failures.

F.6 Closed finite-harmonic controller matrices

For each stored periodic response, the four-component reduced state is represented with $K_G = 5$ harmonics,

$$z(\varphi) = \bar{z} + \sum_{k=1}^{K_G} [a_k \sin(k\varphi) + b_k \cos(k\varphi)]. \quad (\text{F.16})$$

There are $4(1 + 2K_G) = 44$ coefficient perturbations $x \in \mathbb{R}^{44}$. Finite-difference Galerkin projection of the forced 2seg equations gives the local input-output model

$$\begin{aligned} \dot{x} &= A_G x + p_G \Delta\Omega + q_G a + O(\|x, \Delta\Omega, a\|^2), \\ e_\phi &= c_G x + O(\|x\|^2), \quad \Delta\bar{v} = c_v x + O(\|x\|^2), \end{aligned} \quad (\text{F.17})$$

where $a = \Delta A_{\text{exc}}$. Thus $A_G \in \mathbb{R}^{44 \times 44}$, $p_G, q_G \in \mathbb{R}^{44}$, and $c_G, c_v \in \mathbb{R}^{1 \times 44}$.

For this finite-harmonic calculation, the rolling phase estimator is represented locally by a first-order phase channel h , and y denotes the integral contribution:

$$\dot{h} = \omega_c(c_G x - h), \quad \dot{y} = K_I h, \quad \Delta\Omega = s(K_P h + y). \quad (\text{F.18})$$

For the state (x, h, y) , the phase-only matrix used in the eigenvalue calculation is therefore

$$\mathcal{A}_{\text{PLL}} = \begin{bmatrix} A_G & sp_G K_P & sp_G \\ \omega_c c_G & -\omega_c & 0 \\ 0 & K_I & 0 \end{bmatrix}. \quad (\text{F.19})$$

Equivalently, its phase-loop factor has the closed transfer form

$$\det(\lambda I - A_G) [\lambda(\lambda + \omega_c) - \omega_c s(K_P \lambda + K_I) c_G (\lambda I - A_G)^{-1} p_G] = 0. \quad (\text{F.20})$$

At a fixed desired speed, the locally unsaturated perturbation of Equation (F.10) satisfies $\dot{a} = -K_v c_v x$ when $g_\phi = 1$ and no saturation is active. For the state (x, h, y, a) , the matrix used for the two-loop eigenvalues is

$$\mathcal{A}_{\text{PLL+speed}} = \begin{bmatrix} A_G & sp_G K_P & sp_G & q_G \\ \omega_c c_G & -\omega_c & 0 & 0 \\ 0 & K_I & 0 & 0 \\ -K_v c_v & 0 & 0 & 0 \end{bmatrix}. \quad (\text{F.21})$$

The numerical evaluation uses $s = +1$, $\omega_c = \Omega_{\text{exc}}$, $K_P = 0.04$, $K_I = 0.08$, and $K_v = 2.5 \times 10^{-5}$. On the first fourteen ordered amplitude levels, $A_{\text{exc}} = 0.0020496$ through 0.0314870 , the spectral abscissa of $\mathcal{A}_{\text{PLL}}$ lies in $[-1.81848, -1.00100] \times 10^{-3}$; for $\mathcal{A}_{\text{PLL+speed}}$ it lies in $[-8.93731, -2.95865] \times 10^{-4}$. All 28 sampled spectral abscissae are negative. The next amplitude level, $A_{\text{exc}} = 0.0355189$, is not included in this linearized set.

This result concerns the displayed finite-dimensional, unsaturated local models. The largest projected-residual RMS is 2.44015×10^{-2} , whereas the smallest two-loop spectral margin is 2.95865×10^{-4} ; no perturbation bound transfers those values to the discarded harmonics or to the exact nonlinear forced system.

F.7 Requirements for a finite-arrival theorem

A finite-arrival theorem for the stage manager would additionally require an initial set inside a verified response tube, a uniform settling-and-acceptance bound for sufficiently small proposals, a positive minimum accepted progress step, and uniform contraction and guard margins. With the evidence presently available, its verified amplitude interval would end at $A_{\text{exc}} = 0.0314870$ unless a separate local stability argument covered the next numerical endpoint. These conditions are not consequences of the two nominal progressions, so the control result remains local phase-speed regulation supported by sampled finite-harmonic evidence.

Appendix G

Local geometry of continued locomotion families

The words *family*, *branch*, and *manifold* refer to related but different objects. This appendix makes their local mathematical relation explicit. It supports the terminology used in the abstract and conclusion; it also states why a finite numerical collection alone cannot establish one global smooth set.

G.1 Removing the arbitrary time shift

Let z denote a pose-reduced mechanical state and let $\Phi_T(z_0; p)$ be the flow at time T for a fixed model parameter p . A periodic or reduced-periodic orbit has an arbitrary choice of time origin: if $z(t)$ is a solution, then $z(t + \tau)$ describes the same orbit. A section or phase condition removes this redundancy.

Collect the independent return equations and one phase condition in a residual

$$F(u) = 0, \quad u = (z_0, T, \beta) \in \mathbb{R}^N, \quad F : \mathbb{R}^N \rightarrow \mathbb{R}^M, \quad (\text{G.1})$$

where β contains only genuine unknowns beyond the initial state and period: for example a model parameter, a forcing parameter, or an auxiliary unfolding variable. A derived quantity such as energy or mean speed can label the branch after solution. If it is promoted to an unknown, its defining equation must be included among the residual rows. The entries of F must be independent; in particular, a scalar exchange equation that replaces one return row must not also be counted as an additional row.

Proposition G.1 (Regular local branch). *Suppose F is continuously differentiable near a solution u_* and $\text{rank } DF(u_*) = M$. Then the zero set of F is locally a smooth manifold of dimension $N - M$. When $N - M = 1$, it is locally a continuation branch.*

Proof. This is the regular-level-set form of the implicit function theorem. A local choice of M dependent coordinates leaves $N - M$ free coordinates that parameterize the zero set. $\square$

The proposition is conditional on the rank hypothesis. Solver convergence and a visually smooth plotted curve are not substitutes for that hypothesis. Pseudo-arclength continuation can traverse a simple fold of a chosen plotting parameter, but a branch point or other rank loss requires separate analysis (Kuznetsov, 2004).

G.2 From a phase-fixed branch to an invariant family

A phase-fixed branch at one fixed vector field selects one representative on every orbit. Restoring the phase sweeps each representative along its cycle:

$$\mathcal{M}_{\text{loc}} = \{\Phi_\tau(z_0(\alpha); p) : \alpha \in I, \tau \in \mathbb{R}/T(\alpha)\mathbb{Z}\}. \quad (\text{G.2})$$

On a regular embedded portion, this sweep is locally two-dimensional: one coordinate selects the orbit and one selects phase on that orbit. It is invariant because advancing the fixed flow changes only the phase coordinate. If the branch coordinate changes the vector field itself, the branch instead lives in an extended orbit-parameter solution space; its union is not an invariant state-space manifold of any one member of that parameterized family.

This local statement does not imply that every computed branch, every modal sector, or different locomotor models belong to one global manifold. Branches may terminate, meet, bifurcate, or live in state spaces of different dimensions. The phrase *Natural Locomotion Manifold* names the organizing structure; the computations in this thesis provide transported branch candidates and, where transport, rank, and embedding are established, local charts of that structure.




Figure G.1: Local solution geometry and an observable projection. A fold of a plotted coordinate need not destroy regularity of the solution branch, whereas a rank-loss branch point cannot be crossed by invoking the regular-level-set argument. The diagram is conceptual.

G.3 A manifold and its observable image are not the same object

Let h collect reported quantities, for example mean speed, period, energy, phase, or pose increment:

$$h : \mathcal{M}_{\text{loc}} \rightarrow \mathbb{R}^r. \quad (\text{G.3})$$

Even if $\mathcal{M}_{\text{loc}}$ is smooth, its image under h may fold or self-intersect. A fold in a speed–energy plot can therefore be a property of the chosen projection rather than a singularity of the underlying solution set. Conversely, a smooth-looking point cloud does not establish the rank, embeddedness, or global topology of its source.

Figure G.1 contrasts these two distinct phenomena: a regular source branch whose observable image folds, and a genuine rank-loss point that requires separate local charts.

G.4 Natural and resonant locomotion objects

For the conservative calculations, F contains exact return of the reduced state, a phase condition, and the fixed-model equations. Pose reconstruction then selects the transported subset $\Delta g \neq e_G$; only regular embedded local phase sweeps of that subset are called NLM charts. In the two-segment scalar theorem, zero integrated exchange replaces the one remaining crossing speed equation after the other return coordinates have closed. In the three-segment locomotor, full vector return is imposed first and exchange is an endpoint identity or diagnostic; adding it as an independent selector would give a false codimension.

For a periodically forced model, the state may be augmented by forcing phase, or equivalently sampled stroboscopically. For one-to-one synchronization, a residual of the form

$$F_{\text{f}}(z_0, T_{\text{exc}}, A_{\text{exc}}, \dots) = 0 \quad (\text{G.4})$$

defines a synchronized-response solution set $\mathcal{S}_R$. If the independent residual Jacobian has the expected rank, regular portions form local curves or sheets according to the number of free forcing and design coordinates. Pose reconstruction then selects the transported subset

$$\mathcal{M}_R = \{s \in \mathcal{S}_R : \Delta g(s) \neq e_G\}. \quad (\text{G.5})$$

Regular local components of this transported set, continued within the one-to-one resonance construction, are called the *Resonant Locomotion Manifold* (RLM). Here *resonant* identifies the response family organized around the one-to-one resonance; it does not require every member to lie at exact fundamental quadrature. Speed–phase–period diagrams are observable maps, not the definition of smoothness or transport.

Pose reconstruction of all 13,782 recorded responses gives a nonidentity increment at the numerical resolution stated in Section E.5. Recorded closure, unit forcing-clock winding, and the population-wide exact-field local-flow, energy-work, and pose checks therefore support the numerical classification of the synchronized responses as transport-bearing, and hence as members of $\mathcal{M}_R$ at the stated resolution. The retained data do not include the complete collocation Jacobians needed for a population-wide regularity test. Calling a neighbourhood an RLM chart consequently remains conditional on the rank and embedding hypotheses. The mechanics-preserving local conservative-to-forced bridge is a separate calculation and does not prove persistence of the complete conservative family.

G.5 What a numerical regularity check would establish

Where the independent residual and its scaling are reproducible, a useful local check reports the singular values of the scaled Jacobian along a branch. The check must specify:

1. the unknowns and independent residual rows after phase fixing;
2. the scaling used for variables with different units or magnitudes;
3. the smallest singular value or numerical rank along regular segments;
4. the locations excluded as branch points, modal limits, or missing data.

Such a calculation would support local-chart language on the examined segments. It would still not prove a single global NLM or RLM across different models, modal sectors, bifurcations, or parameterizations.

Bibliography

- Gaëtan Abeloos, Felix Müller, Emre Ferhatoglu, Matthias Scheel, Christophe Collette, Gaëtan Kerschen, Matthew R. W. Brake, Paolo Tiso, Ludovic Renson, and Malte Krack. A consistency analysis of phase-locked-loop testing and control-based continuation for a geometrically nonlinear frictional system. *Mechanical Systems and Signal Processing*, 170: 108820, 2022. doi: 10.1016/j.ymssp.2022.108820.
- Boye K. Ahlborn, Robert W. Blake, and William M. Megill. Frequency tuning in animal locomotion. *Zoology*, 109(1):43–53, 2006. doi: 10.1016/j.zool.2005.11.001.
- Alin Albu-Schäffer and Cosimo Della Santina. A review on nonlinear modes in conservative mechanical systems. *Annual Reviews in Control*, 50:49–71, 2020. doi: 10.1016/j.arcontrol.2020.10.002.
- Vladimir I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer, 2 edition, 1989. doi: 10.1007/978-1-4757-2063-1.
- David A. W. Barton. Control-based continuation: Bifurcation and stability analysis for physical experiments. *Mechanical Systems and Signal Processing*, 84:54–64, 2017. doi: 10.1016/j.ymssp.2015.12.039.
- Filip Bjelonic, Arne Sachtler, Alin Albu-Schäffer, and Cosimo Della Santina. Experimental closed-loop excitation of nonlinear normal modes on an elastic industrial robot. *IEEE Robotics and Automation Letters*, 7(2):1689–1696, 2022. doi: 10.1109/LRA.2022.3141156.
- Anthony M. Bloch. *Nonholonomic Mechanics and Control*. Springer, 2 edition, 2015. doi: 10.1007/978-1-4939-3017-3.
- Frédéric Boyer and Ayman Belkhir. Reduced locomotion dynamics with passive internal DoFs: Application to nonholonomic and soft robotics. *IEEE Transactions on Robotics*, 30(3):578–592, 2014. doi: 10.1109/TRO.2013.2294733.
- Bernard Brogliato. *Nonsmooth Mechanics: Models, Dynamics and Control*. Springer, 3 edition, 2016. doi: 10.1007/978-3-319-28664-8.
- Davide Calzolari, Cosimo Della Santina, and Alin Albu-Schäffer. Exciting families of passive gaits in an elastic quadruped via natural motion manifold control. *The International Journal of Robotics Research*, page 02783649251347305, 2025. doi: 10.1177/02783649251347305.
- Carlos Canudas-de Wit, Bernard Espiau, and Claudio Urrea. Orbital stabilization of underactuated mechanical systems. *IFAC Proceedings Volumes*, 35(1):527–532, 2002. doi: 10.3182/20020721-6-ES-1901.00900.
- Mattia Cenedese and George Haller. How do conservative backbone curves perturb into forced responses? a Melnikov function analysis. *Proceedings of the Royal Society A*, 476(2234):20190494, 2020a. doi: 10.1098/rspa.2019.0494.
- Mattia Cenedese and George Haller. Stability of forced–damped response in mechanical systems from a Melnikov analysis. *Chaos*, 30(8):083103, 2020b. doi: 10.1063/5.0012480.
- Robert Clewley. Hybrid models and biological model reduction with PyDSTool. *PLOS Computational Biology*, 8(8):e1002628, 2012. doi: 10.1371/journal.pcbi.1002628.
- Steven H. Collins, Andy Ruina, Russ Tedrake, and Martijn Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085, 2005. doi: 10.1126/science.1107799.
- Alessandro De Luca, Giuseppe Oriolo, and Paolo Robuffo Giordano. Image-based visual servoing schemes for nonholonomic mobile manipulators. *Robotica*, 25(2):131–145, 2007. doi: 10.1017/S0263574706003262.
- Cosimo Della Santina and Alin Albu-Schäffer. Exciting efficient oscillations in nonlinear mechanical systems through eigenmanifold stabilization. *IEEE Control Systems Letters*, 5(6):1916–1921, 2021. doi: 10.1109/LCSYS.2020.3048228.

- Vincent Denis, Marc Jossic, Christophe Giraud-Audine, Benoît Chomette, Alexandre Renault, and Olivier Thomas. Identification of nonlinear modes using phase-locked-loop experimental continuation and normal form. *Mechanical Systems and Signal Processing*, 106:430–452, 2018. doi: 10.1016/j.ymssp.2018.01.014.
- Eusebius J. Doedel, Alan R. Champneys, Fabio Dercole, Thomas F. Fairgrieve, Yuri A. Kuznetsov, Bart E. Oldeman, Randy C. Paffenroth, Björn Sandstede, Xianjun Wang, and Chenghai Zhang. *AUTO-07P: Continuation and Bifurcation Software for Ordinary Differential Equations*, 2007.
- Jaap Eldering and Henry O. Jacobs. The role of symmetry and dissipation in biolocomotion. *SIAM Journal on Applied Dynamical Systems*, 15(1):24–59, 2016. doi: 10.1137/140970914.
- Adrien Escande, Nicolas Mansard, and Pierre-Brice Wieber. Hierarchical quadratic programming: Fast online humanoid-robot motion generation. *The International Journal of Robotics Research*, 33(7):1006–1028, 2014. doi: 10.1177/0278364914521306.
- Francesco Fasso, Simone Passarella, and Marta Zoppello. Control of locomotion systems and dynamics in relative periodic orbits. *Journal of Geometric Mechanics*, 12(3):395–420, 2020. doi: 10.3934/jgm.2020022.
- Vitaliy Fedonyuk and Phanindra Tallapragada. The dynamics of a chaplygin sleigh with an elastic internal rotor. *Regular and Chaotic Dynamics*, 24(1):114–126, 2019. doi: 10.1134/S1560354719010076.
- Zhenyu Gan, Yevgeniy Yesilevskiy, Petr Zaytsev, and C. David Remy. All common bipedal gaits emerge from a single passive model. *Journal of the Royal Society Interface*, 15(146):20180455, 2018. doi: 10.1098/rsif.2018.0455.
- Ernst Hairer, Syvert P. Nørsett, and Gerhard Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*. Springer, 2 edition, 1993. doi: 10.1007/978-3-540-78862-1.
- George Haller and Sten Ponsioen. Nonlinear normal modes and spectral submanifolds: Existence, uniqueness and use in model reduction. *Nonlinear Dynamics*, 86(3):1493–1534, 2016. doi: 10.1007/s11071-016-2974-z.
- Ross L. Hatton and Howie Choset. Geometric motion planning: The local connection, Stokes’ theorem, and the importance of coordinate choice. *The International Journal of Robotics Research*, 30(8):988–1014, 2011. doi: 10.1177/0278364910394392.
- Navvab Kashiri, Andy Abate, Sabrina J. Abram, Alin Albu-Schäffer, Patrick J. Clary, Monica Daley, Salman Faraji, Raphael Furnemont, Manolo Garabini, Hartmut Geyer, Alena M. Grabowski, Jonathan Hurst, Jörn Malzahn, Glenn Mathijssen, C. David Remy, Wesley Roosting, Mohammad Shahbazi, Surabhi N. Simha, Jae-Bok Song, Nils Smit-Anseeuw, Stefano Stramigioli, Bram Vanderborght, Yevgeniy Yesilevskiy, and Nikos Tsagarakis. An overview on principles for energy efficient robot locomotion. *Frontiers in Robotics and AI*, 5:129, 2018. doi: 10.3389/frobt.2018.00129.
- Scott D. Kelly and Richard M. Murray. Geometric phases and robotic locomotion. *Journal of Robotic Systems*, 12(6):417–431, 1995. doi: 10.1002/rob.4620120607.
- Gaëtan Kerschen, Maxime Peeters, Jean-Claude Golinval, and Alexander F. Vakakis. Nonlinear normal modes, part i: A useful framework for the structural dynamicist. *Mechanical Systems and Signal Processing*, 23(1):170–194, 2009. doi: 10.1016/j.ymssp.2008.04.002.
- Robert J. Kuether, Ludovic Renson, Thibaut Detroux, Chiara Grappasonni, Gaëtan Kerschen, and Matthew S. Allen. Nonlinear normal modes, modal interactions and isolated resonance curves. *Journal of Sound and Vibration*, 351:299–310, 2015. doi: 10.1016/j.jsv.2015.04.035.
- Yuri A. Kuznetsov. *Elements of Applied Bifurcation Theory*. Springer, 3 edition, 2004. doi: 10.1007/978-1-4757-3978-7.
- Nicolas Mansard, Oussama Khatib, and Abderrahmane Kheddar. A unified approach to integrate unilateral constraints in the stack of tasks. *IEEE Transactions on Robotics*, 25(3):670–685, 2009. doi: 10.1109/TRO.2009.2020345.
- Tad McGeer. Passive dynamic walking. *The International Journal of Robotics Research*, 9(2):62–82, 1990. doi: 10.1177/027836499000900206.
- Mirado Mortel, Luc Jaulin, Lionel Lapierre, and Simon Rohou. Natural locomotion: Principle and method, 2026. arXiv:2605.28254.
- Jim P. Ostrowski and Joel W. Burdick. The geometric mechanics of undulatory robotic locomotion. *The International Journal of Robotics Research*, 17(7):683–701, 1998. doi: 10.1177/027836499801700701.

- Maxime Peeters, Régis Vigié, Guillaume Sérandour, Gaëtan Kerschen, and Jean-Claude Golinval. Nonlinear normal modes, part ii: Toward a practical computation using numerical continuation techniques. *Mechanical Systems and Signal Processing*, 23(1):195–216, 2009. doi: 10.1016/j.ymssp.2008.04.003.
- Maxime Peeters, Gaëtan Kerschen, and Jean-Claude Golinval. Dynamic testing of nonlinear vibrating structures using nonlinear normal modes. *Journal of Sound and Vibration*, 330(3):486–509, 2011. doi: 10.1016/j.jsv.2010.08.028.
- Simon Peter and Remco I. Leine. Excitation power quantities in phase resonance testing of nonlinear systems with phase-locked-loop excitation. *Mechanical Systems and Signal Processing*, 96:139–158, 2017. doi: 10.1016/j.ymssp.2017.04.011.
- Simon Peter, Robin Riethmüller, and Remco I. Leine. Tracking of backbone curves of nonlinear systems using phase-locked-loops. In *Nonlinear Dynamics, Volume 1*, pages 107–120. Springer, 2016. doi: 10.1007/978-3-319-29739-2_11.
- Hari Krishna Hari Prasad, Ross L. Hatton, and Kaushik Jayaram. Geometric mechanics of contact-switching systems. *IEEE Robotics and Automation Letters*, 8(12):8343–8349, 2023. doi: 10.1109/LRA.2023.3327930.
- Maximilian Raff, Nelson Rosa, and C. David Remy. Connecting gaits in energetically conservative legged systems. *IEEE Robotics and Automation Letters*, 7(3):8407–8414, 2022. doi: 10.1109/LRA.2022.3186500.
- Mirado Rajaomarasata, Luc Jaulin, Lionel Lapierre, and Simon Rohou. Natural efficient gaits from nonholonomic locomotion nonlinear normal mode (NL-NNM): The pendrivencar case. *Mechatronics*, 110:103366, 2025a. doi: 10.1016/j.mechatronics.2025.103366.
- Mirado Rajaomarasata, Luc Jaulin, Lionel Lapierre, and Simon Rohou. Energy-efficient nonholonomic fish robot: Nonlinear forced oscillations. In *OCEANS 2025 Brest*, pages 1–7. IEEE, 2025b. doi: 10.1109/OCEANS58557.2025.11104676.
- Sophie Ramananarivo, Ramiro Godoy-Diana, and Benjamin Thiria. Rather than resonance, flapping wing flyers may play on aerodynamics to improve performance. *Proceedings of the National Academy of Sciences*, 108(15):5964–5969, 2011. doi: 10.1073/pnas.1017910108.
- C. David Remy. *Optimal Exploitation of Natural Dynamics in Legged Locomotion*. PhD thesis, ETH Zurich, 2011.
- Ludovic Renson, Alicia Gonzalez-Buelga, David A. W. Barton, and Simon A. Neild. Robust identification of backbone curves using control-based continuation. *Journal of Sound and Vibration*, 367:145–158, 2016. doi: 10.1016/j.jsv.2015.12.035.
- Robert M. Rosenberg. On nonlinear vibrations of systems with many degrees of freedom. In *Advances in Applied Mechanics*, volume 9, pages 155–242. Elsevier, 1966. doi: 10.1016/S0065-2156(08)70008-5.
- Layale Saab, Oscar E. Ramos, François Keith, Nicolas Mansard, Philippe Souères, and Jean-Yves Fourquet. Dynamic whole-body motion generation under rigid contacts and other unilateral constraints. *IEEE Transactions on Robotics*, 29(2):346–362, 2013. doi: 10.1109/TRO.2012.2234351.
- Matthias Scheel, Simon Peter, Remco I. Leine, and Malte Krack. A phase resonance approach for modal testing of structures with nonlinear dissipation. *Journal of Sound and Vibration*, 435:56–73, 2018. doi: 10.1016/j.jsv.2018.07.010.
- Johan K. Schönebaum, Farbod Alijani, and Giuseppe Radaelli. Review on mobile robots that exploit resonance. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 235(24):7907–7924, 2021. doi: 10.1177/09544062211036923.
- Steven W. Shaw and Christophe Pierre. Normal modes for non-linear vibratory systems. *Journal of Sound and Vibration*, 164(1):85–124, 1993. doi: 10.1006/jsvi.1993.1198.
- Phanindra Tallapragada. A swimming robot with an internal rotor as a nonholonomic system. In *2015 American Control Conference*, pages 657–662. IEEE, 2015. doi: 10.1109/ACC.2015.7170810.
- Min Tang, Christian Stéphan, and Marc Böswald. Phase resonance method for nonlinear mechanical structures with phase locked loop control. In *Proceedings of ISMA 2020*, pages 1805–1818, 2020.
- Eric D. Tytell, Chia-Yu Hsu, and Lisa J. Fauci. The role of mechanical resonance in the neural control of swimming in fishes. *Zoology*, 117(1):48–56, 2014. doi: 10.1016/j.zool.2013.10.011.
- Eric R. Westervelt, Jessy W. Grizzle, and Daniel E. Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE Transactions on Automatic Control*, 48(1):42–56, 2003. doi: 10.1109/TAC.2002.806653.