Control-aided construction of basins of attraction of nonlinear dynamical systems Université de Liège

Auteur : Benlachhab, Reda
Promoteur(s) : Raze, Ghislain
Faculté : Faculté des Sciences appliquées
Diplôme : Master en ingénieur civil en aérospatiale, à finalité spécialisée en "aerospace engineering"
Année académique : 2025-2026
URI/URL : http://hdl.handle.net/2268.2/26173

UNIVERSITY OF LIÈGE · FACULTY OF APPLIED SCIENCES

MASTER'S THESIS

Control-aided construction of basins of attraction of nonlinear dynamical systems

AUTHOR

BENLACHHAB Reda

SUPERVISOR

RAZE Ghislain

ACADEMIC YEAR

2025-2026

SUBMITTED IN FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF

Master of Science in Aerospace Engineering

EXAMINATION COMMITTEE

PRESIDENT

Prof. Gaëtan Kerschen

University of Liège

SUPERVISOR

Dr. Ghislain Raze

University of Liège

MEMBERS

Prof. Olivier Brüls

University of Liège

Prof. Loïc Salles

University of Liège

Abstract

Nonlinear mechanical systems subject to harmonic excitation can exhibit multistability, where qualitatively different steady-state responses coexist under identical operating conditions. The set of initial conditions converging to a given attractor, known as its basin of attraction (BoA), determines which response is observed in practice and how robust it is to perturbations. While numerical methods for basin reconstruction are well established, their experimental counterpart remains comparatively undevelopped: mapping a basin requires placing a physical system at many precisely specified states in the phase plane, simultaneously in displacement and velocity, which no existing experimental technique achieves in a systematic and automated way.

This thesis proposes and demonstrates a control-aided method for the experimental reconstruction of basins of attraction of a Duffing oscillator. A proportional-derivative controller tracks a reference trajectory to drive the system from rest to the neighbourhood of any prescribed target state. A systematic reference design strategy is introduced to ensure energetic consistency between the reference and the true system dynamics across the full phase plane. Control gains are scheduled automatically from the local backbone frequency, and a stroboscopic Poincaré criterion monitors convergence after the controller is released. A trajectory-based backfilling strategy further reduces the number of required experimental runs by propagating attractor labels along confirmed free-evolution paths.

The method is validated numerically on four Duffing configurations of increasing complexity and implemented in real time on a dSPACE MicroLabBox controlling an electronic analogue Duffing oscillator. Experimental basin maps agree closely with numerical references across all configurations, including cases with coexisting subharmonic attractors.

Keywords: Nonlinear dynamics, Basin of attraction, Control-based testing, Duffing oscillator, Reference trajectory, Poincaré map

Acknowledgements

I would like to thank my parents and my brother Rayan for their constant support and for always believing in me.

To my friend and brother Younès, thank you for the conversations and the advice that went well beyond this thesis.

To my classmates who became friends along the way, they will recognise themselves, thank you for making these years what they were.

I would also like to thank the members of the Space Structures and Systems Laboratory (S3L) for their availability and their advice whenever it was needed.

Finally, Ghislain, thank you. Your availability and kindness made this experience far richer than I could have hoped for. I am proud to have been your first supervised student.

Contents

List of Figures

List of Tables

1 Introduction

1.1 Context and Motivation

Vibrations are an omnipresent phenomenon in aerospace engineering. Every flight structure, from the wing of a commercial aircraft to the solar panels of a satellite, is subject to dynamic excitation arising from aerodynamic loads, engine operation, and structural interactions. Understanding how a structure responds to these excitations is therefore a central concern of structural dynamics, and it underpins decisions made at every stage of the design process, from material selection to certification testing.

Classical structural dynamics rests on the assumption of linearity. Under this assumption the equations of motion are linear in the displacements and velocities, the principle of superposition holds, and the particular response at any forcing frequency is uniquely determined by the excitation. The framework is mature, computationally tractable, and supported by a rich body of experimental modal analysis techniques. For many practical structures and loading levels it provides an adequate description of the dynamics.

The increasing use of advanced composite materials, thin flexible panels, and geometrically nonlinear joints in modern aerospace structures is, however, pushing the boundaries of this linear paradigm. Lightweight wings undergo large deflections during gusts and manoeuvres; thermally loaded panels develop post-buckled equilibria; vibration energy harvesters are deliberately designed to exploit nonlinear stiffness characteristics. In all of these configurations the restoring forces depend nonlinearly on displacement, and the linear approximation fails to capture the essential dynamics. The tools developed for linear systems arrive at their limits, and new methods are required.

One of the most distinctive consequences of nonlinearity is the possibility of multistability: under identical forcing conditions the system can settle into qualitatively different long-term responses depending on where it started. A post-buckled beam may vibrate about either of its two buckled equilibria; a hardening oscillator may sustain either a low-amplitude or a high-amplitude resonant orbit at the same excitation frequency. The geometry of the initial conditions that lead to a given attractor is called the basin of attraction, and it carries crucial information about the dynamics. In particular, it determines which response is actually observed in operation, how robust that response is to external perturbations, and whether a dangerous large-amplitude attractor can be reached accidentally from a nominally safe operating point. Knowledge of the basin structure is therefore not only of academic interest: it is a prerequisite for reliable design, safe operation, and targeted control of nonlinear mechanical systems.

Despite its importance, and the existence of a number of mature numerical tools for its computation, the experimental measurement of a basin of attraction remains a largely unsolved problem. The difficulty is fundamental: a basin map requires the system to be placed, sequentially, at each point of a dense grid of initial conditions in the phase space, each specified simultaneously in displacement and velocity. Standard experimental methods provide no mechanism for imposing arbitrary initial conditions on a running system in a repeatable and automated way. State-of-the-art methods, reviewed in Chapter 3 of this thesis, only offer a partial solution to these issues.

1.2 Objectives

This thesis proposes a method that revolves around two stages to address these shortcomings, which is the central contribution of this work. In Stage 1, a proportional-derivative controller tracks an analytically prescribed reference orbit to drive the system from a fixed resting state at the origin to the neighbourhood of any prescribed target point in the phase plane. In Stage 2, the controller is switched off and the system evolves freely until a stroboscopic Poincaré criterion declares convergence to a periodic attractor. Repeating this two-stage procedure over a dense grid of target points maps the complete basin of attraction.

The specific objectives of the work are as follows.

1. Design a control-aided method for the experimental determination of basins of attraction of nonlinear oscillators and provide a tuning methodology for the control parameters.

2. Implement the algorithm in real time on a dSPACE MicroLabBox controlling an electronic analogue Duffing oscillator circuit, and demonstrate automated cell-by-cell experimental basin reconstruction.

3. Validate the two-stage algorithm numerically and experimentally across qualitatively distinct configurations of a Duffing oscillator featuring various flavours of nonlinear dynamics.

4. Investigate the impact of the choice of the reference orbit on the performance of the proposed method.

The method is developed and demonstrated for single-degree-of-freedom systems, for which the basin of attraction is a two-dimensional subset of the phase plane and can be mapped at practical grid densities. Extension to systems with more degrees of freedom would require either a targeted sampling strategy or a dimensionality reduction to the most influential phase-space directions; this is not pursued here. The electronic Duffing circuit used for experimental validation is an analogue model of a mechanical oscillator, not a physical mechanical structure. The circuit provides clean, repeatable, and well-characterised nonlinear dynamics at practical time scales, making it an ideal validation platform. The portability of the algorithm to physical mechanical systems, where state measurement involves accelerometers and displacement sensors rather than voltage probes, is discussed in the conclusion but not experimentally demonstrated herein.

1.3 Thesis Outline

Chapter 2 establishes the theoretical foundations, reviewing the linear oscillator as a baseline and the Duffing oscillator as the central model, deriving the exact Jacobi elliptic solutions of the conservative system, and defining the basin of attraction formally. Chapter 3 reviews the state of the art in basin reconstruction, identifies the gap that motivates the controlled approach, and develops the full two-stage algorithm including the reference design, gain scheduling, cell-entry criterion, and Poincaré convergence detector. Chapter 4 presents the numerical validation across the four Duffing configurations, including the frequency sweep and parameter uncertainty analyses. Chapter 5 presents the experimental platform and the experimental basin reconstructions for all four cases, with quantitative comparison against the numerical ground truth. Chapter 6 draws conclusions, states the original contributions of the work, and identifies directions for future research.

2 Theoretical Background

The dynamic behaviour of mechanical structures is governed by Newton's second law. For a single-degree-of-freedom (SDOF) oscillator of mass m, subject to a restoring force $g(x, \dot{x})$ and a harmonic excitation $F_0 \cos(\Omega t)$, this reads

$$m\ddot{x} + g(x, \dot{x}) = F_0 \cos(\Omega t). \tag{1}$$

Provided g is globally Lipschitz continuous in its arguments, the Picard-Lindelöf theorem guarantees existence and uniqueness of the solution for given initial conditions [11]. Section 2.1 treats the linear oscillator as a reference baseline, establishing three key properties of linear systems, namely (i) forced solutions are unique, (ii) they obey a superposition principle, and (iii) they are globally asymptotically stable, all of which will be violated in the nonlinear case. Section 2.2 extends the analysis to the Duffing oscillator. Section 2.3 formalises the basin of attraction. Section 2.4 briefly addresses the MDOF(Multi-Degree-of-Freedom) extension. Throughout, the state vector $\mathbf{y} = (x, \dot{x})^\top \in \mathbb{R}^2$ fully describes the system at any instant.

2.1 Linear Oscillator

2.1.1 Equation of Motion

Expanding g in a Taylor series about the equilibrium $x = 0$ and retaining only first-order terms gives $g \approx kx + c\dot{x}$, where $k = \partial g/\partial x\big|_0$ is the linear stiffness and $c = \partial g/\partial \dot{x}\big|_0$ the linear damping coefficient. Equation (1) becomes

$$m\ddot{x} + c\dot{x} + kx = F_0 \cos(\Omega t). \tag{2}$$

Introducing the natural frequency $\omega_0 = \sqrt{k/m}$ and the damping ratio $\zeta = c/(2m\omega_0)$ yields the normalised form

$$\ddot{x} + 2\zeta\omega_0\dot{x} + \omega_0^2 x = \frac{F_0}{m}\cos(\Omega t). \tag{3}$$

Because (3) is linear, its general solution obeys the principle of superposition: it is the sum of the homogeneous solution (the transient free response, decaying exponentially for $\zeta > 0$) and a particular solution (the steady-state forced response). In state-space form with $\mathbf{y} = (x, \dot{x})^\top$:

$$\dot{\mathbf{y}} = A\mathbf{y} + \mathbf{b}\cos(\Omega t), \qquad A = \begin{bmatrix} 0 & 1 \\ -\omega_0^2 & -2\zeta\omega_0 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 0 \\ F_0/m \end{bmatrix}. \tag{4}$$

2.1.2 Energy

The mechanical energy of the conservative part of the system is

$$E(x, \dot{x}) = \frac{1}{2}m\dot{x}^2 + \frac{1}{2}kx^2. \tag{5}$$

Differentiating along trajectories of (2):

$$\dot{E} = m\dot{x}\ddot{x} + kx\dot{x} = \dot{x}(m\ddot{x} + kx) = F_0\dot{x}\cos(\Omega t) - c\dot{x}^2. \tag{6}$$

The analysis of this scalar quantity proves useful to understand the global dynamics of both unforced and forced cases, as shown below.

2.1.3 Unforced Dynamics

When $F_0 = 0$, a constant solution $\mathbf{y}^*$ of the vector field satisfies $f(\mathbf{y}^*) = \mathbf{0}$, giving the unique equilibrium $x^* = 0$, $\dot{x}^* = 0$, which is simultaneously the only zero of the restoring force and the global minimum of E.

In the damped case ($c > 0$), equation (6) with $F_0 = 0$ gives $\dot{E} = -c\dot{x}^2 \le 0$: energy is non-increasing and strictly decreasing whenever $\dot{x} \neq 0$. LaSalle's invariance principle [9] then guarantees convergence to the largest invariant subset of $\{\dot{E} = 0\}$, i.e. $\dot{x} \equiv 0$, which forces $x \equiv 0$. The origin is therefore a global attractor: every trajectory converges to it, irrespective of the initial condition, and the long-term behaviour is uniquely determined by the parameters alone.

In the undamped case ($c = 0$), equation (6) gives $\dot{E} = 0$: the system is conservative and the equilibrium at the origin becomes a stable centre surrounded by iso-energy ellipses. The exact solution of equation (3) is

$$x(t) = A\cos(\omega_0 t + \phi_0), \tag{7}$$

with amplitude A and phase ϕ_0 set by initial conditions. Every trajectory oscillates at the same frequency ω_0, regardless of amplitude, a property called isochrony that will be lost in the nonlinear case. Figure 1 illustrates the contrast between the closed elliptical orbit of the undamped system and the inward-spiralling trajectory of the damped one.

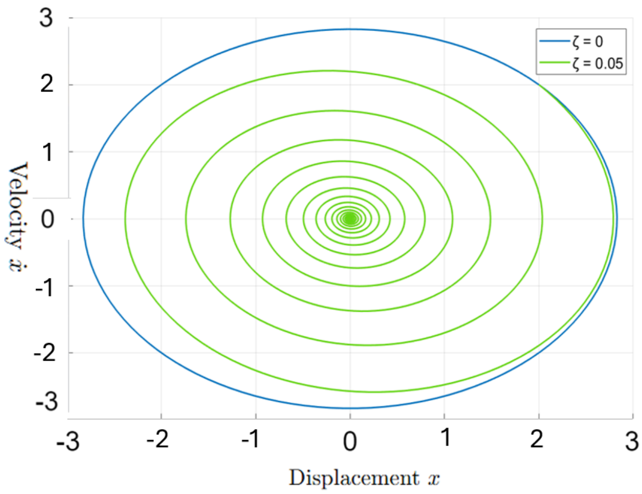

Figure 1: Phase portrait of a linear oscillator. The closed ellipse is a constant-energy contour of the undamped system; the spiral is the trajectory of the damped system, crossing iso-energy ellipses inward as energy dissipates.

2.1.4 Forced Dynamics

Under harmonic excitation the transient decays and the unique steady-state solution of equation (3) is $x(t) = X(\Omega)\cos(\Omega t + \phi)$, where ϕ is a phase and the amplitude $X(\Omega)$ is given by

$$X(\Omega) = \frac{F_0/k}{\sqrt{(1-r^2)^2 + (2\zeta r)^2}}, \qquad r = \frac{\Omega}{\omega_0}. \tag{8}$$

The steady-state amplitude and phase are shown in Figure 3. Near the frequency ratio $r = 1$, the characteristics of resonance are observed through a peak in the amplitude and a phase close to 90. At $r = 1$, the power balance $\dot{E} = 0$ holds exactly in steady state, meaning the power injected by the forcing is entirely absorbed by damping.

Three properties of (8) serve as the reference baseline for the comparison with nonlinear behaviour. First, $X(\Omega)$ is single-valued: at any forcing frequency there is exactly one steady-state amplitude. Second, X scales proportionally with F_0 so superposition holds. Third, for $c > 0$ the steady-state is globally asymptotically stable: the long-term response does not depend on the initial condition. None of these three properties survives when a nonlinear restoring force is introduced.

Figure 2 illustrates the decay of the transient and the emergence of the steady-state response.

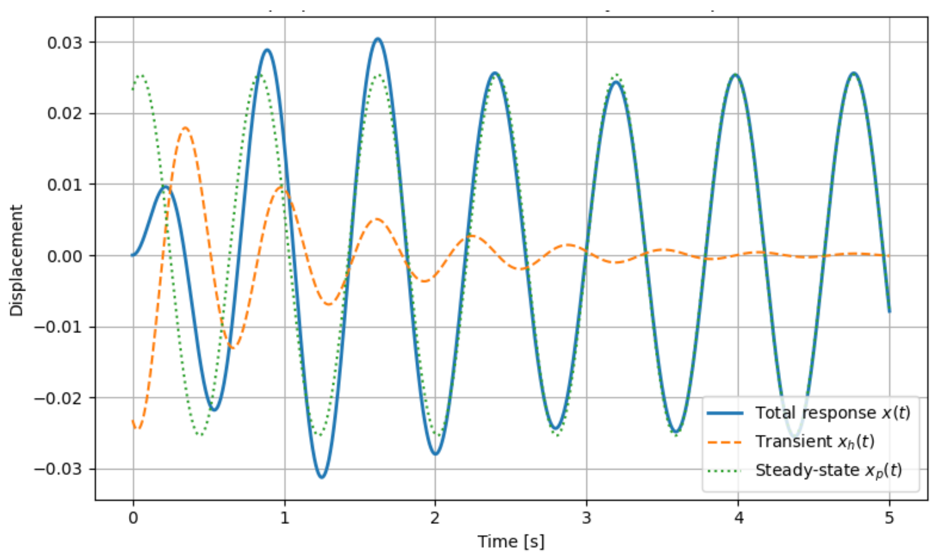

Figure 2: Time response of a linear oscillator to harmonic excitation. The dashed curve shows the decaying transient; the solid curve shows the total response converging to the steady-state sinusoid.

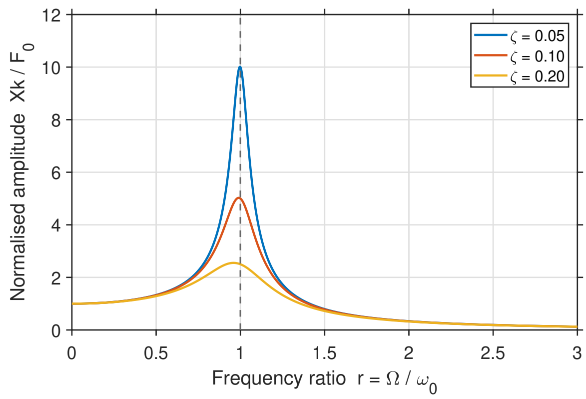

(a) Amplitude X/X_{st} as a function of the frequency ratio $r = \Omega/\omega_0$ for three values of ζ. Each curve is single-valued and scales linearly with F_0.

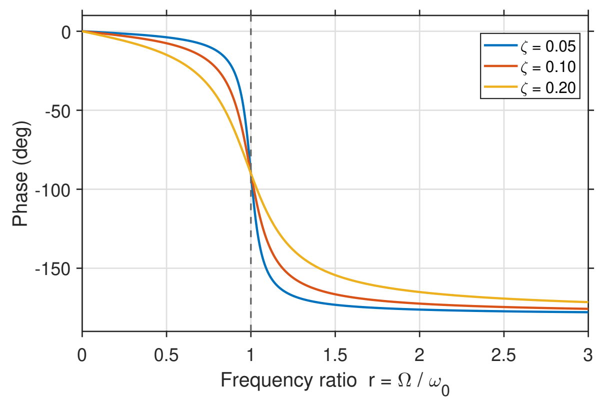

(b) Phase lag ϕ versus r. The transition through $-\pi/2$ at resonance ($r = 1$) is independent of F_0.

Figure 3: Frequency response of a linear SDOF oscillator for three damping ratios ζ.

2.2 The Duffing Oscillator

2.2.1 Motivation and Equation of Motion

Many engineering systems deviate measurably from linear behaviour even at moderate amplitudes: post-buckled beams, clamped thin plates, cable-stayed structures, and vibration energy harvesters all exhibit restoring forces that depend nonlinearly on displacement [15].

When the restoring force is odd-symmetric, $g(-x, -\dot{x}) = -g(x, \dot{x})$, the Taylor expansion contains only odd powers of x. Retaining the linear and cubic terms gives the Duffing oscillator

$$\ddot{x} + \delta\dot{x} + \alpha x + \beta x^3 = \gamma \cos(\omega t), \tag{9}$$

where $\delta \geq 0$ is the viscous damping coefficient, α the linear stiffness, β the cubic coefficient, γ the forcing amplitude, and ω the forcing frequency. Related asymmetric models (Helmholtz, Helmholtz-Duffing) are not considered here [8].

Two parameter regimes are central to this work. In the monostable hardening regime ($\alpha > 0$, $\beta > 0$) the restoring force stiffens with amplitude. In the bistable double-well regime ($\alpha < 0$, $\beta > 0$) the linear stiffness is negative, creating a double-well potential with two stable equilibria. This configuration arises in post-buckled beams [22] and in vibration energy harvesters [13]. Because equation (9) is nonlinear, the principle of superposition no longer holds, which is responsible for the violation of all three baseline properties identified in Section 2.1.

2.2.2 Energy and Potential Landscape

The mechanical energy is

$$E(x, \dot{x}) = \frac{1}{2}\dot{x}^2 + V(x), \qquad V(x) = \frac{\alpha}{2}x^2 + \frac{\beta}{4}x^4, \tag{10}$$

where V is the potential associated with the conservative part of the restoring force. For $\alpha > 0$, V has a single minimum at $x = 0$. For $\alpha < 0$, V has a local maximum at $x = 0$ and two symmetric minima at $x_\pm = \pm\sqrt{-\alpha/\beta}$; the separatrix energy $E_s = -\alpha^2/(4\beta)$ is the energy of the saddle at the origin and defines the boundary between in-well and cross-well motion. As shown below, the shape of this potential greatly influences both the unforced and forced dynamics.

Differentiating along trajectories of (9):

$$\dot{E} = \dot{x}(\ddot{x} + V'(x)) = \dot{x}\gamma\cos(\omega t) - \delta\dot{x}^2. \tag{11}$$

This mirrors the structure of (6): the first term represents power injected by the external forcing, and the second represents energy dissipated by damping. The geometry of the potential $V(x)$, shown in Figure 4, is what qualitatively distinguishes the Duffing oscillator from its linear counterpart.

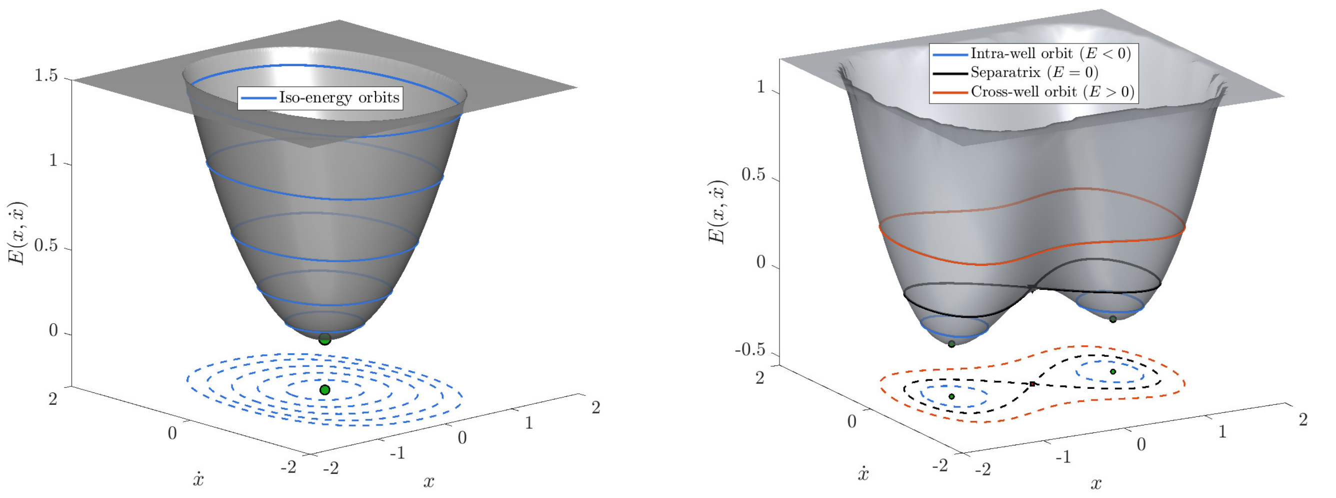

(a) Monostable potential ($\alpha = 1$, $\beta = 1$). All iso-energy contours are closed curves encircling the single equilibrium at the origin.

(b) Bistable double-well potential ($\alpha = -1$, $\beta = 1$). Contours below E_s (dashed) close around one well; those above encircle both wells.

Figure 4: Potential energy $V(x)$ and energy surface $E(x, \dot{x}) = \frac{1}{2}\dot{x}^2 + V(x)$ with iso-energy contours projected onto the phase plane.

2.2.3 Unforced Dynamics

When $\gamma = 0$, constant solutions of the unforced system satisfy $\alpha x^* + \beta x^{*3} = 0$, giving $x^* = 0$ always and $x_\pm = \pm\sqrt{-\alpha/\beta}$ when $\alpha < 0$. Local stability can be assessed with the Jacobian at each equilibrium point, given by

$$J(x^*) = \begin{bmatrix} 0 & 1 \\ -(\alpha + 3\beta x^{*2}) & -\delta \end{bmatrix}. \tag{12}$$

From its eigenvalues, it can be shown that the local stiffness $\alpha + 3\beta x^{*2}$ determines the stability type: a positive value yields a stable focus (for $\delta > 0$); a negative value yields an unstable saddle.

In the damped case ($\delta > 0$), equation (11) with $\gamma = 0$ gives $\dot{E} = -\delta \dot{x}^2 \leq 0$. LaSalle's invariance principle and this inequality guarantee convergence to equilibrium points. For the monostable system the origin is the unique global attractor. For the bistable system, which equilibrium is reached depends entirely on the initial condition, as illustrated in Figure 5.

In the undamped case ($\delta = 0$), $\dot{E} = 0$ and the system evolves along iso-energy curves. The orbital period is now amplitude-dependent and the isochrony is lost. The exact solutions are Jacobi elliptic functions, given here for a case with zero initial velocity and initial displacement A [8]. Three orbit families exist for the bistable case: in-well dn orbits ($E < E_s$), cross-well cn orbits ($E > E_s$), and the non-periodic separatrix ($E = E_s$, a homoclinic connection). For the in-well case:

$$x(t) = \pm A\,\mathrm{dn}\big(\Omega_{\mathrm{dn}}(t + t_0) \mid m''\big)\,, \qquad \Omega_{\mathrm{dn}} = A\sqrt{\tfrac{\beta}{2}}, \qquad m'' = \frac{-2\alpha}{\beta A^2} \in (0, 1), \tag{13}$$

where dn is the delta amplitude function and m'' is its elliptic modulus. For the monostable case ($\alpha > 0$) all orbits are cn-type:

$$x(t) = A\,\mathrm{cn}\big(\Omega(t + t_0) \mid m'\big)\,, \qquad \Omega = \sqrt{\alpha + \beta A^2}, \qquad m' = \frac{\beta A^2}{2(\alpha + \beta A^2)} \in \big[0, \tfrac{1}{2}\big), \tag{14}$$

where cn is the elliptic cosine function and m' is its elliptic modulus. As $\beta \to 0$ or $A \to 0$, $m' \to 0$ and $\mathrm{cn}(u \mid 0) = \cos u$, recovering the sinusoidal linear solution (7). The backbone curve $\Omega(A) = \sqrt{\alpha + \beta A^2}$ shows that the natural frequency increases with amplitude for $\beta > 0$. The elliptic modulus m' quantifies the departure of the orbit from a pure sinusoid: for $m' \approx 0$ the motion is nearly harmonic, while for $m' \to \frac{1}{2}$ it becomes strongly anharmonic. This parameter plays a central role in the reference design of Chapter 3.4.

The trajectories can be visualised in projections onto the phase plane $(x, \dot{x})$ in Figure 5.

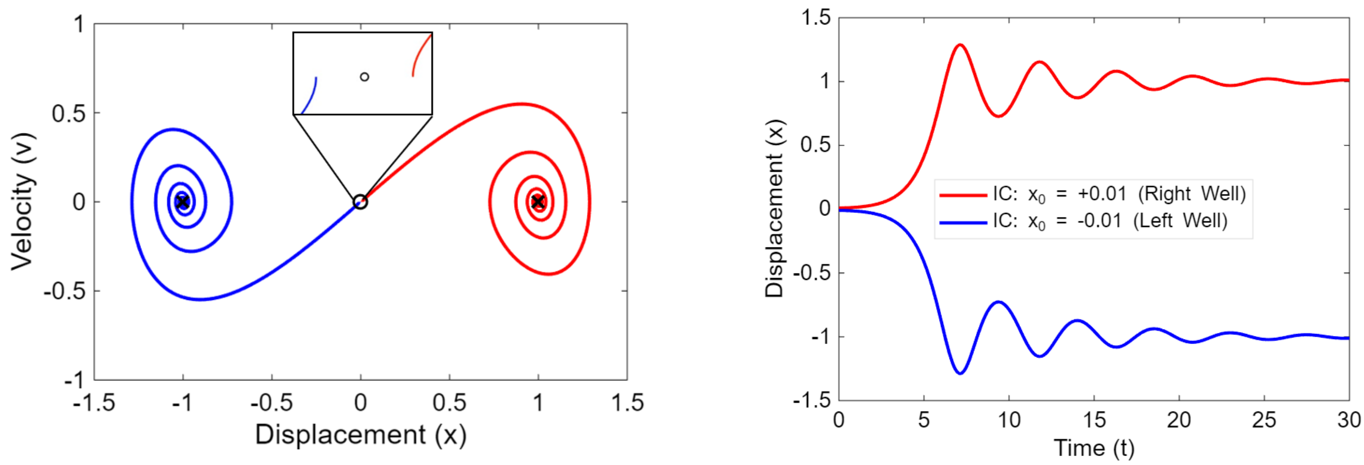

(a) Phase portrait. Two trajectories from close initial conditions converge to opposite equilibria.

(b) Corresponding time responses. Despite close starting positions the long-term outcomes are opposite.

Figure 5: Unforced bistable Duffing oscillator ($\gamma = 0$, $\delta = 0.3$, $\alpha = -1$, $\beta = 1$). Two initial conditions symmetric about the origin converge to opposite equilibria, illustrating that the long-term behaviour depends critically on the starting point.

2.2.4 Forced Dynamics: Nonlinear Frequency Response

Under harmonic excitation the relevant long-term solutions may be periodic orbits locked to the forcing period $T_f = 2\pi/\omega$ (period-1 attractors) and, for appropriate parameters, subharmonic attractors of period pT_f for integer $p \geq 2$. The FRC bends toward higher frequencies for hardening systems ($\beta > 0$), reflecting the amplitude-dependent backbone $\Omega(A)$. At resonance, trajectories in the phase space are close to the exact undamped solution when forcing and damping are small. Over a finite frequency band the curve becomes multi-valued: two stable and one unstable periodic orbit coexist at the same forcing frequency (Figure 6). The transitions occur at fold bifurcations (saddle-node bifurcations of periodic orbits), bounding the bistable band. Stability can be assessed with Floquet theory [20]. The middle branch is typically unstable and cannot be captured reliably through direct time integration; instead, it must be followed using continuation techniques [18, 12, 15]. In contrast, the two outer branches correspond to stable solutions and represent the two coexisting attractors of the system.

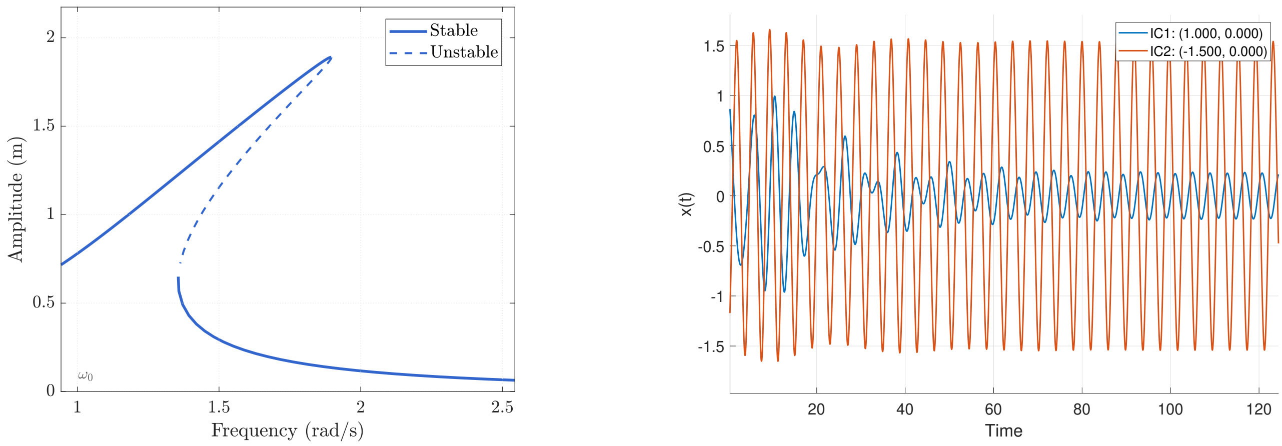

(a) FRC of the hardening Duffing oscillator ($\alpha = 1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0.35$). Solid: stable branches. Dashed: unstable branch. The dash-dot vertical line marks the operating frequency used in Case 1.

(b) Coexisting period-1 attractors at the same forcing frequency $\omega = 1.6$ rad/s, reached from two different initial conditions. All three baseline properties of the linear FRC are violated.

Figure 6: Nonlinear FRC and attractor coexistence of the hardening Duffing oscillator ($\alpha = 1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0.35$).

For sufficiently large forcing amplitudes, isolas appear: closed FRC branches entirely disconnected from the main resonance (Figure 7). An isola encloses stable large-amplitude periodic orbits that are generally inaccessible from rest by any quasi-static frequency sweep [4]. They can only be reached by placing the system directly within their basin of attraction; without controlled initial-condition placement, an isola can go entirely undetected experimentally.

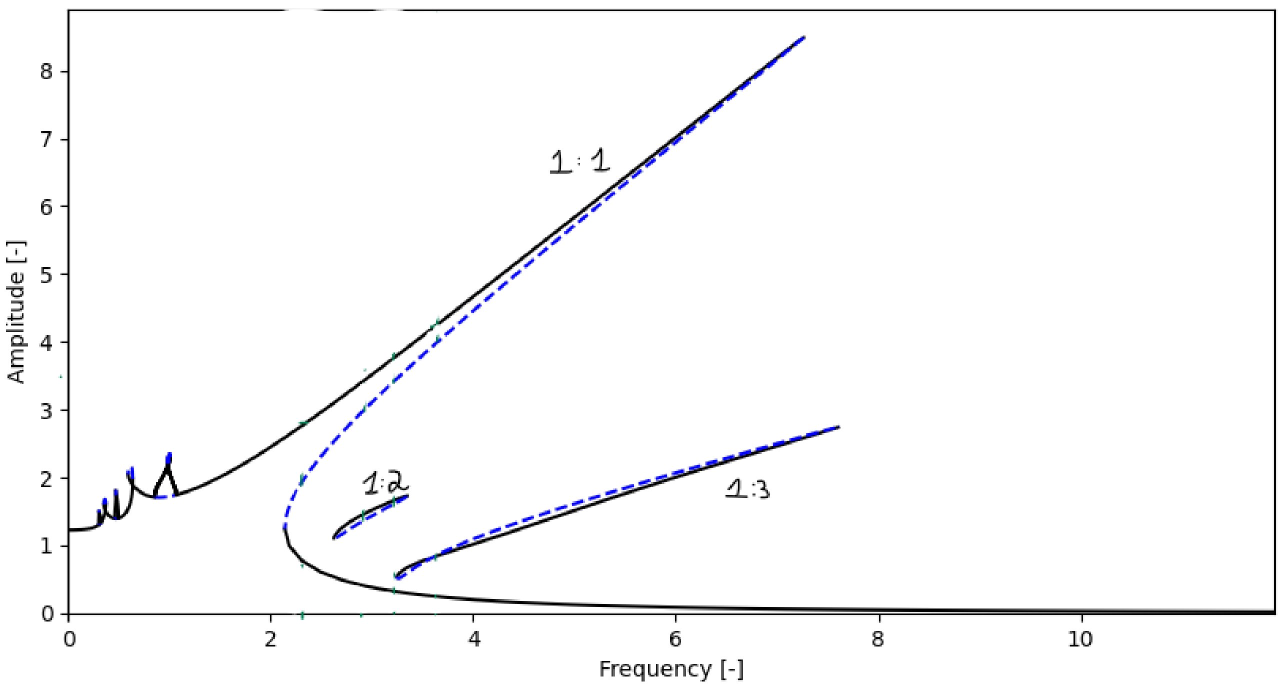

Figure 7: Frequency-response curve exhibiting isolas obtained using numerical continuation [2] for $\alpha = 1$, $\beta = 1$, $\delta = 0.05$, $\gamma = 3$. The disconnected closed branches correspond to subharmonic resonances.

2.3 Global Analysis and the Basin of Attraction

The coexistence of multiple stable attractors means that the long-term behaviour of a nonlinear system is not uniquely determined by its parameters but also depends critically on the initial condition $\mathbf{y}_0 \in \mathbb{R}^n$, as discussed in standard nonlinear dynamics treatments [20]. The basin of attraction of a bounded attractor $\mathbf{a}_k(t)$ is defined as the set of all initial conditions whose trajectories converge asymptotically to that attractor:

$$\mathcal{B}(\mathbf{a}_k) = \left\{ \mathbf{y}_0 \in \mathbb{R}^n : \lim_{t\to\infty} \left\| \mathbf{y}(t; \mathbf{y}_0) - \mathbf{a}_k(t) \right\| = 0 \right\}. \tag{15}$$

Here $\mathbf{a}_k$ may represent a fixed point, periodic orbit, or a more general invariant set, and $\mathbf{y}(t; \mathbf{y}_0)$ denotes the unique trajectory starting from $\mathbf{y}_0$.

The phase space is thus partitioned into disjoint basins of attraction,

$$\mathbb{R}^n = \mathcal{B}(\mathbf{a}_1) \cup \cdots \cup \mathcal{B}(\mathbf{a}_K) \cup \mathcal{S}, \tag{16}$$

where $\mathcal{S}$ is a set of measure zero containing initial conditions that do not converge to any of the K attractors. In practice, the structure and extent of these basins are not accessible through local analysis and must instead be characterised using global numerical sampling of initial conditions, a standard approach in nonlinear dynamics [19].

Standard analysis tools each fail to provide basin information. Local linearisation (Jacobian, eigenvalues) characterises stability only in a small neighbourhood of each attractor. Harmonic balance computes the frequency response and detects bifurcation points but does not assign basin membership. Numerical continuation traces solution branches and bifurcations with respect to parameters but remains fundamentally local to a single branch at a time. Global analysis consists of integrating trajectories from many initial conditions and recording their asymptotic fate. It is therefore required to directly reconstruct $\mathcal{B}$. Its importance is threefold: it provides a complete picture of the dynamics over the full domain of interest; it identifies the full set of attractors, including isola attractors that are invisible to frequency sweeps; and it quantifies the relative importance of each attractor, since a large basin implies the corresponding response is likely to be observed in practice, while a small basin signals high sensitivity to initial conditions.

Figure 8 shows the basin map for the unforced bistable Duffing oscillator of Figure 5. Initial conditions that lie in the blue and red areas eventually converge to the right and left equilibria, respectively. The basin boundary coincides with the stable manifold of the saddle at the origin and is smooth and symmetric because the vector field shares the same symmetry. This explains the results of Figure 5: while close, the two initial conditions considered therein lie on opposite sides of the basin boundary and therefore converge to distinct attractors.

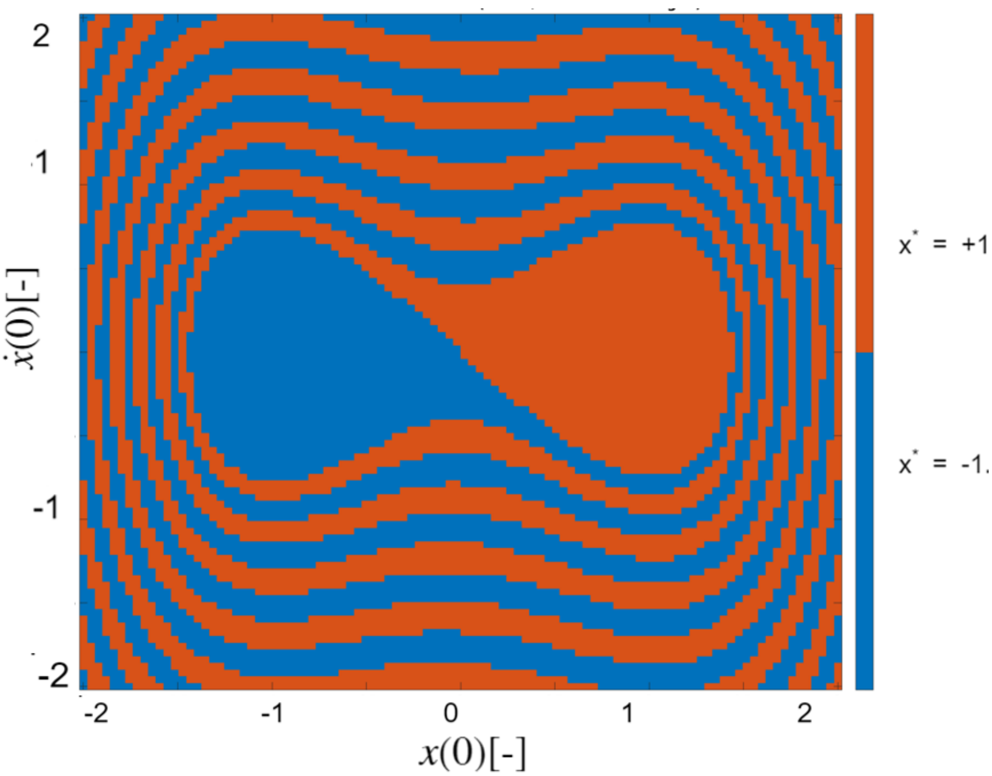

Figure 8: Basin of attraction for the unforced bistable Duffing oscillator ($\alpha = -1$, $\beta = 1$, $\delta = 0.3$). Blue: initial conditions converging to $x_+ = +1$. Red: initial conditions converging to $x_- = -1$.

Figure 9 shows the basin map for the hardening Duffing oscillator ($\alpha = 1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0.35$) under forcing at a frequency $\omega = 1.6$ rad/s inside the bistable band. Similarly to the unforced bistable case, the two coloured regions depict which among the low- and high-amplitude attractors shown in Figure 6a will be reached for given initial conditions. Unlike the unforced case, the basin boundary develops a complex interleaved structure near the stable manifold of the unstable periodic orbit, reflecting high sensitivity to initial conditions in this regime.

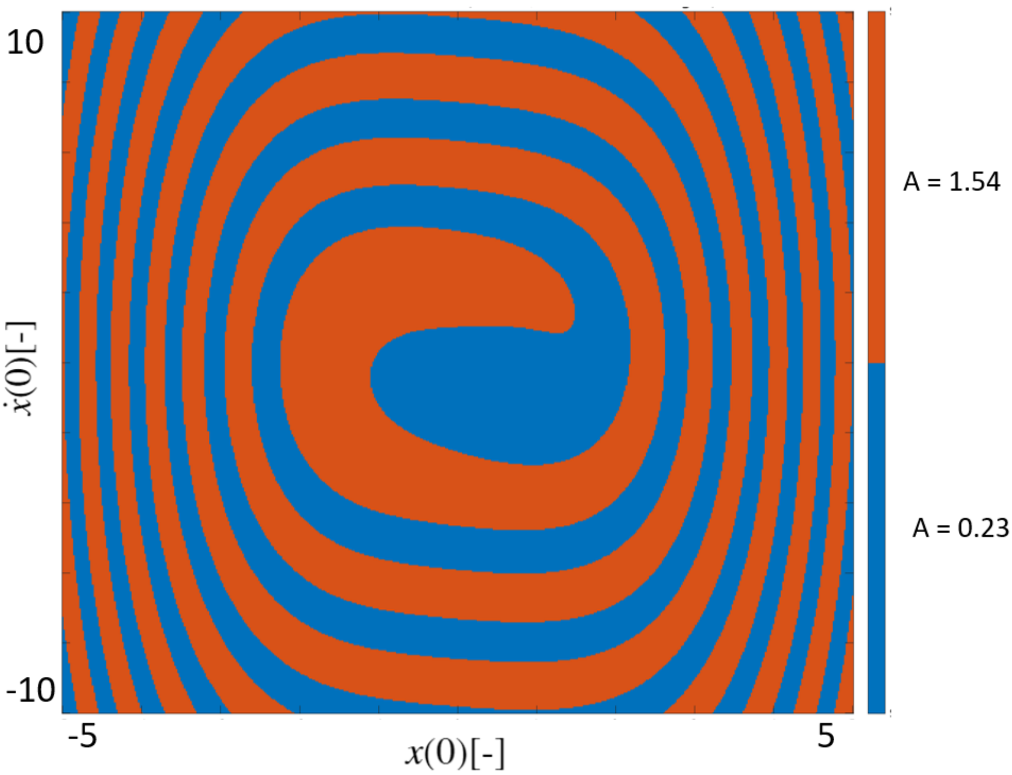

Figure 9: Basin of attraction for the forced hardening Duffing oscillator ($\alpha = 1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0.35$) at forcing frequency $\omega = 1.6$ rad/s inside the bistable band.

2.4 Extension to MDOF Systems

The concepts above extend naturally to systems with $n > 2$ degrees of freedom governed by

$$M\ddot{\mathbf{q}} + C\dot{\mathbf{q}} + K\mathbf{q} + \mathbf{g}_{\mathrm{nl}}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{F}(t). \tag{17}$$

For linear MDOF systems, the principle of superposition holds and the response can be decomposed into independent modal contributions. In contrast, nonlinear MDOF systems exhibit fundamentally different behaviour: superposition breaks down, and modal interactions arise through nonlinear coupling. A key mechanism is internal resonance, which occurs when natural frequencies satisfy a rational relation $\omega_i/\omega_j \in \mathbb{Q}$, leading to sustained inter-modal energy transfer and the emergence of solution branches that cannot be captured by single-mode descriptions [10]. This loss of modal independence and the resulting nonlinear coupling are central features of nonlinear normal mode theory, where it is shown that classical linear modal decomposition fails and that new solution branches may arise due to mode interaction [7].

Basin analysis in $\mathbb{R}^n$ becomes increasingly challenging as the number of degrees of freedom increases. For $n \geq 4$, the full basin structure is no longer directly visualisable and its numerical determination becomes computationally expensive, requiring targeted sampling strategies rather than full-grid exploration [1]. More recent developments also highlight that identifying basins in high-dimensional systems requires specialised data-driven or trajectory-based discrimination techniques to overcome the curse of dimensionality. This thesis therefore focuses on the SDOF case ($n = 2$), for which the phase plane provides a complete and interpretable representation of the global dynamics.

3 Construction of Basins of Attraction

Having established in Section 2.3 that the basin of attraction is the fundamental object governing global dynamics, this section addresses the practical question of how to measure it in an experimental setting. Three existing families of methods are reviewed in Section 3.1 and the gap motivating the present work is identified. The proposed method is then presented as an overview in Section 3.2, and developed in detail through Sections 3.3-3.6.

3.1 State of the Art

Several numerical methods exist for the global analysis of dynamical systems and are relatively mature. Sequential numerical integration, parallel integration over dense grids, and cell-mapping methods [6] all allow systematic reconstruction of basins of attraction from a mathematical model. The situation is altogether different in an experimental context. In principle, identifying a mathematical model for an experimental set-up is possible, after which numerical methods could be applied to the model. In practice, however, nonlinear system identification remains a significant challenge, and it is desirable, when possible, to avoid relying on a model entirely. The three families of experimental methods reviewed below allow this, at least to some extent.

3.1.1 Stochastic Interrogation

Stochastic interrogation [21] periodically switches between a nominal excitation and a random perturbation. The perturbation has the consequence that the system evolves through random states, which are seen as random initial conditions once the nominal excitation is switched back on. This effectively releases the system from a large ensemble of randomly chosen initial conditions, and the asymptotic outcome of each trajectory can then be recorded.

The principal advantage of this approach is its experimental simplicity: no explicit model and no active control are required. Each realisation is obtained by passively releasing the system from an uncontrolled perturbation, and the resulting long-term behaviour is classified according to the observed attractor.

However, the method is intrinsically limited by uncertainty in the initial release conditions. The true starting point cannot be imposed exactly, since the release mechanism introduces an uncontrolled perturbation whose magnitude is difficult to quantify. As a consequence, the reconstructed basin is a statistically blurred representation of the true basin structure; fine features of the basin boundary cannot be resolved below the standard deviation of the initial release error. A further disadvantage is that certain regions of the state space are not well explored, as the random perturbations do not provide uniform coverage; this can be observed in Figure 10.

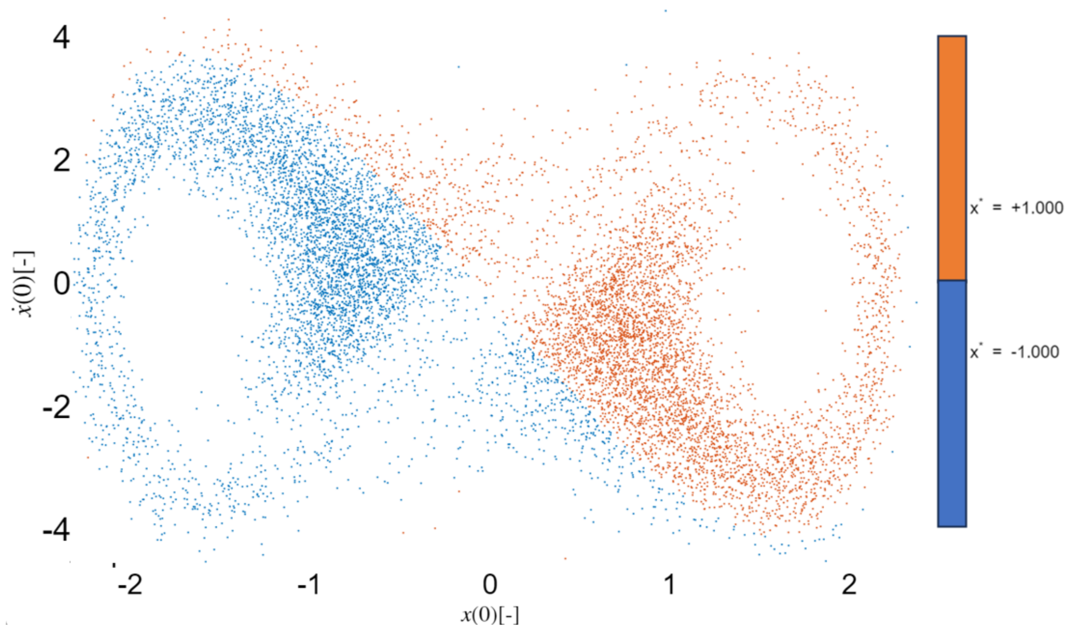

Figure 10: Stochastic interrogation basin for the unforced bistable case using 10^6 random initial conditions.

3.1.2 Parametric Modification

A second family of methods infers basin structure indirectly by monitoring bifurcations as a system parameter is varied quasi-statically. The moment at which the system jumps from one branch to another during a frequency sweep is an experimental observation of a fold bifurcation, and the fold frequencies bound the bistable band. Within this band the ratio of basin areas can, in principle, be estimated from the probability of observing each attractor after a perturbation [17].

An alternative within this family consists in switching the system parameters to steer the system onto different attractors, then switching back to the nominal configuration in order to explore multiple initial conditions. However, it is not always possible to apply such parameter switching, and the achievable precision on the resulting initial conditions is difficult to control.

These methods are experimentally accessible but provide only indirect and coarse information about basin geometry. They locate fold bifurcations and associated erosion curves, but they do not resolve the cell-by-cell basin boundary and cannot detect disconnected basins associated with isola attractors.

3.1.3 Direct Initial-Condition Imposition

The third family attempts to place the system at a prescribed state $(x^*, \dot{x}^*)$ by active mechanical interventions, using mechanical stops, impulsive forcing, or short-duration control inputs. Compared with stochastic interrogation, direct imposition improves repeatability and allows systematic grid coverage. The fundamental difficulty is that position and velocity must be imposed simultaneously: a mechanical stop constrains only x, and imposing $\dot{x}$ simultaneously would require an additional actuator or timing strategy. In published studies the achievable accuracy is typically of the order of the grid cell size [6].

None of the three families combines precise simultaneous imposition of both state components, systematic dense grid coverage, and a mechanism to reach disconnected or energetically remote basin regions. The strategy developed in the remainder of this section is designed to address all three requirements for the Duffing oscillator.

3.2 Control-Aided Construction of Experimental Basins of Attraction

A method that revolves around two stages is proposed to address these shortcomings and is the central contribution of this thesis. The purpose of the first stage is to reach a given initial condition $(x_*, \dot{x}_*)$, and that of the second stage is to determine the attractor associated with this initial condition. The procedure can be repeated for each initial condition, thereby mapping the state space and allowing for the construction of basins of attraction in an experiment with barely any post-processing. An overview of the method is shown in Figure 11.

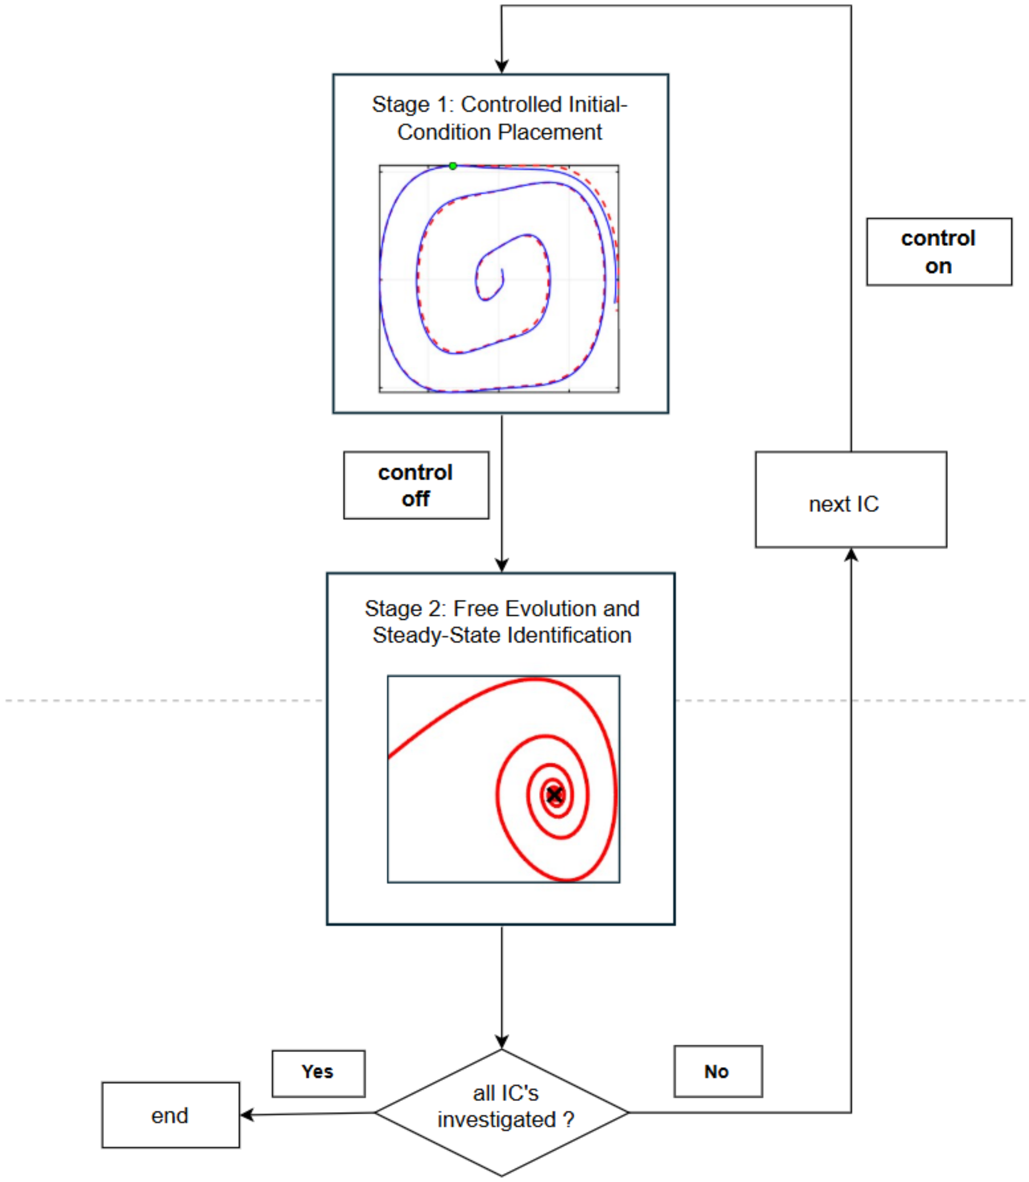

Figure 11: Overview of the proposed two-stage method for the construction of basins of attraction.

In the first stage, the system is driven through feedback control to follow a trajectory that intersects the initial condition of interest. Elliptic target trajectories can be used by default, but more suitable trajectories can be selected based on the characteristics of the system under test. Appropriate control gains are selected depending on the trajectory. The target trajectory can also be smoothly adapted at the start of the procedure to avoid excessive control effort during transients. Finally, a method to detect when the system reaches the desired state is proposed.

In the second stage, the control is turned off and the system evolves in its operating conditions. Autonomous systems are thus no longer forced, or the external excitation signal of a forced system is reset at release time. The target state from which the system is released effectively acts as an initial condition, which eventually reveals its associated attractor. This requires an adequate sampling method and an associated steady-state detection strategy.

In the sequel, basic but important features of feedback control are recalled in Section 3.3. The first and second phases of the method are presented in detail in Sections 3.4 and 3.5, respectively.

3.3 Control Theory Reminder

The controlled basin reconstruction strategy relies on a feedback controller that actively drives the system to a prescribed state. This section provides the control-theoretic background necessary to follow the subsequent development.

3.3.1 Open-Loop and Closed-Loop Control

Open-loop control computes the actuating signal entirely from the desired output, without observing what the plant actually does. Because no information flows back from the plant to the controller, any mismatch between the design model and the true plant accumulates without correction. Closed-loop (feedback) control overcomes this limitation by continuously measuring the plant output and feeding it back to the controller, which then acts on the error: the signed difference between the desired reference and the measured output. The error signal is

$$e(t) = r(t) - y(t), \tag{18}$$

where $r(t)$ is the reference and $y(t)$ the measured output. The self-correcting mechanism makes the system robust to disturbances and model uncertainty: any deviation from the desired behaviour is detected and counteracted automatically.

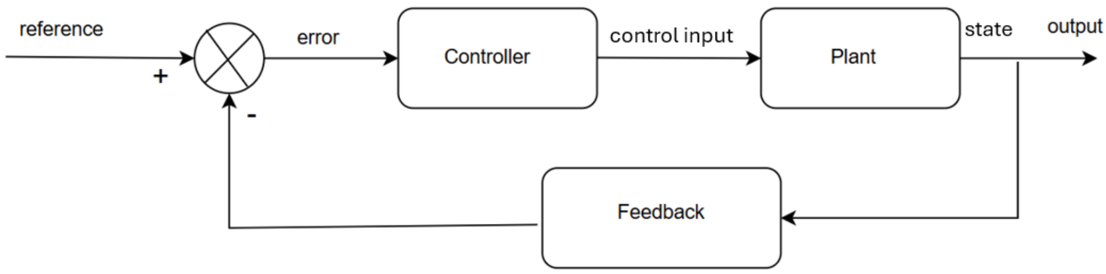

Figure 12: Standard closed-loop block diagram.

3.3.2 The PID Controller

The proportional-integral-derivative (PID) controller synthesises a control input u based on the error signal according to

$$u(t) = K_p\, e(t) + K_i \int_0^t e(\tau)\,\mathrm{d}\tau + K_d\, \dot{e}(t), \tag{19}$$

where K_p, K_i, and K_d are the proportional, integral, and derivative gains. As an illustrative example, the regulation of a second-order mass-spring-damper system is considered hereafter. The open-loop plant is governed by

$$m\ddot{x}(t) + c\dot{x}(t) + kx(t) = u(t), \tag{20}$$

where $x(t)$ represents the output position, $m = 1$ is the mass, $c = 0.2$ the damping coefficient, and $k = 1$ the structural stiffness. The objective is to make the step response of the closed-loop system as close to the reference $r(t) = 1$ as possible in a fast and robust way. The system is simulated numerically over $t \in [0, 10]$ s using a standard Runge-Kutta solver (`ode45`) under zero initial conditions. The discussion is centred on performance and control effort, but not on stability and robustness, for brevity.

Proportional action. The proportional term $K_p e(t)$ produces a corrective effort proportional to the instantaneous error. As shown in Figure 13, increasing K_p accelerates the response but reduces the steady-state offset at the cost of increased overshoot and a larger initial control action magnitude.

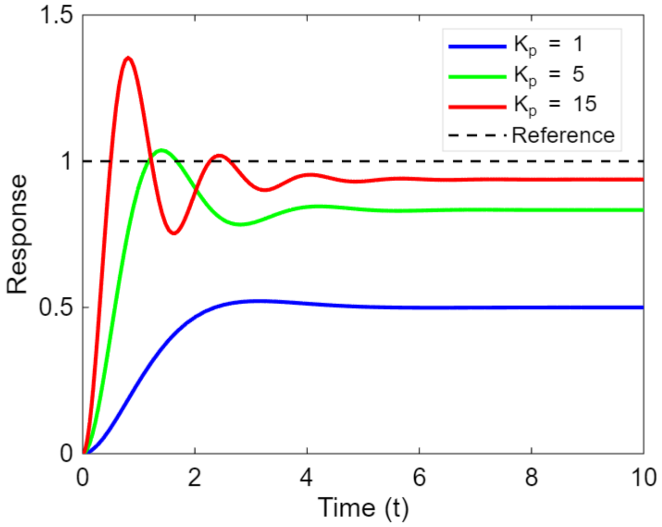

(a) System tracking response $x(t)$.

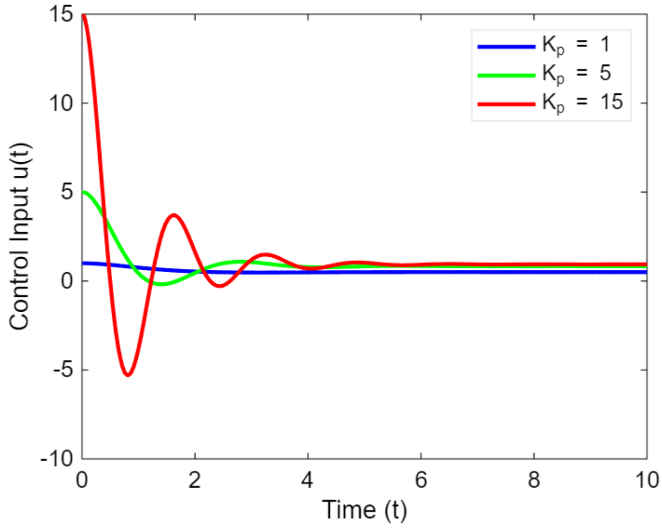

(b) Control input signal $u(t)$.

Figure 13: Step response and control effort of a P controller for three values of K_p. Larger K_p reduces steady-state offset but increases overshoot and spikes transient control effort.

Integral action. The integral term $K_i \int_0^t e(\tau)\,\mathrm{d}\tau$ accumulates the error over time and eliminates the steady-state offset left by a purely proportional controller under constant disturbance. However, integral action can cause oscillations and destabilisation if the gain is too large, as illustrated in Figure 14.

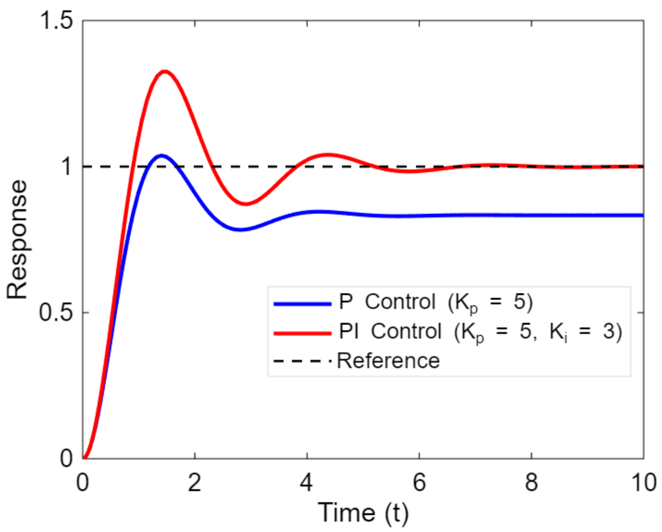

(a) System tracking response $x(t)$.

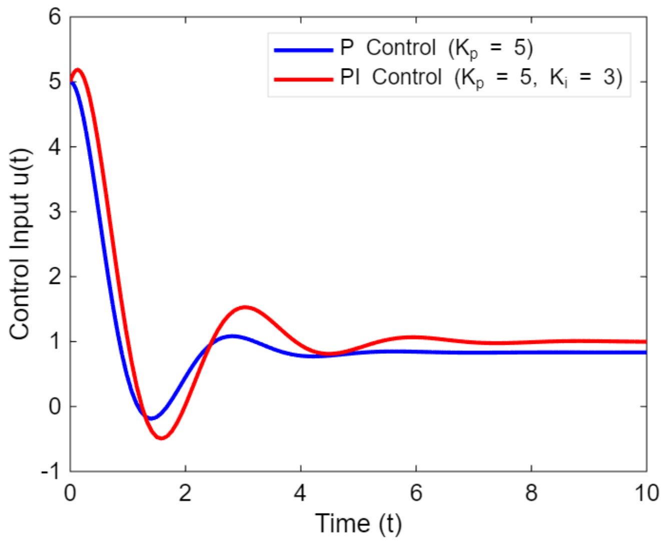

(b) Control input signal $u(t)$.

Figure 14: Step response and control effort of a PI controller. The integral action eliminates residual offset by maintaining a constant steady-state control effort, at the cost of added oscillations.

Derivative action. The derivative term $K_d\dot{e}(t)$ reacts to the rate of change of the error, providing a predictive damping mechanism that reduces overshoot and improves transient performance. It can, however, slow the system down if too large, and it amplifies measurement noise. Figure 15 illustrates this behaviour.

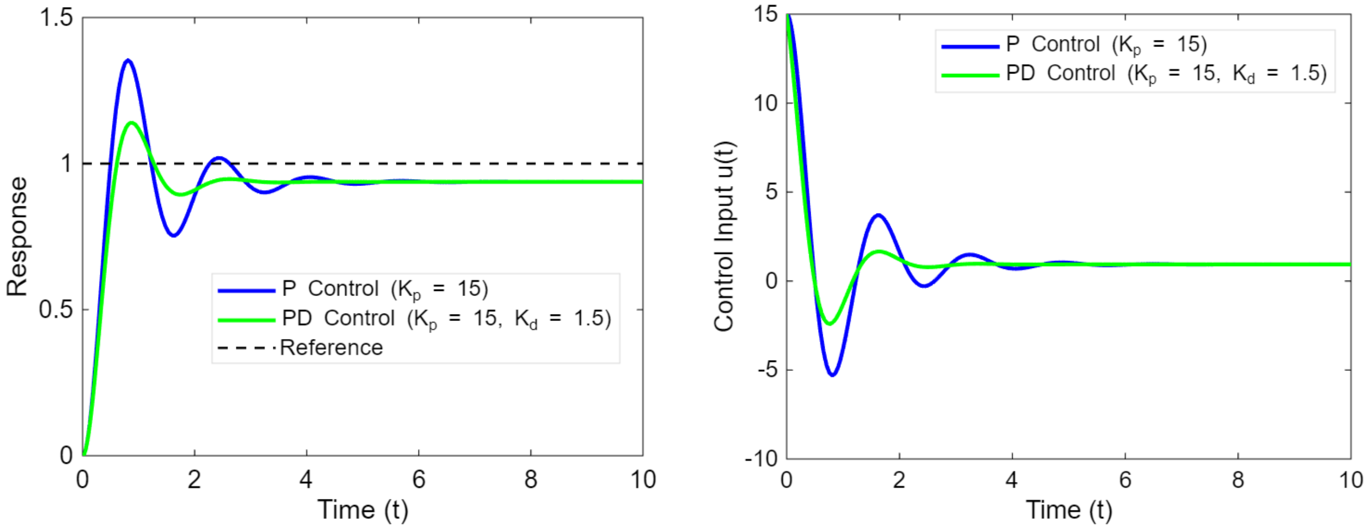

(a) System tracking response $x(t)$. (b) Control input signal $u(t)$.

Figure 15: Step response and control effort of a PD controller. Derivative action actively counters fast changes in error, providing a damping feature that curtails overshoot.

Overall, choosing large gains improves tracking performance but increases the control input magnitude.

3.3.3 Choice of PD Control for Trajectory Tracking

The control objective of this work is trajectory tracking: the system must pass through a prescribed state $(x^*, \dot{x}^*)$ specifying both position and velocity simultaneously. Regarding the three PID terms in the context of this objective, the following choices are made.

The proportional gain K_p is important for good tracking performance and is therefore retained. The integral gain K_i is less important here because the reference is not a step but a sinusoidal or elliptic trajectory. Since the drawback of potential destabilisation outweighs the advantage of static error elimination for the cases considered herein, integral action is not retained. The derivative gain K_d is important for stabilisation and is therefore retained. Since in the experimental set-up it is possible to directly measure the velocity, the derivative term is implemented through direct velocity feedback. This removes the need to evaluate derivatives numerically and thus mitigates the noise issue usually associated with the derivative term of the PID.

The resulting PD controller is

$$u(t) = K_p\big(x_r(t) - x(t)\big) + K_d\big(\dot{x}_r(t) - \dot{x}(t)\big), \tag{21}$$

where $x_r(t)$ and $\dot{x}_r(t)$ are the reference position and velocity, and $K_p, K_d > 0$. The proportional term corrects the position error; the derivative term acts on the velocity error. Together they drive the system toward a moving reference orbit. Figure 16 shows the corresponding closed-loop block diagram.

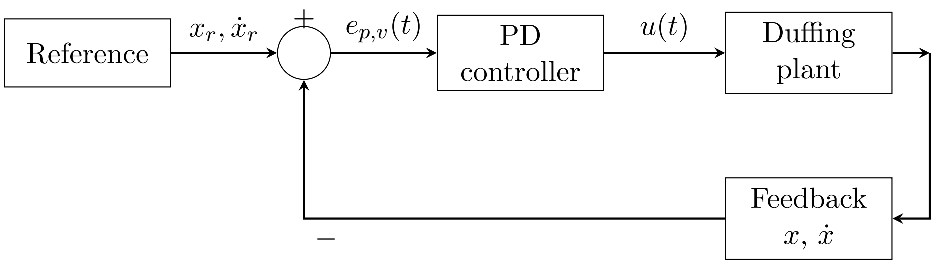

Figure 16: Closed-loop block diagram of the PD-controlled Duffing plant. The controller computes $u(t)$ from the position and velocity errors with respect to the reference trajectory. Full state feedback closes the loop.

The design of the reference trajectory $x_r(t)$ and the gain pair (K_p, K_d) are the two principal degrees of freedom, addressed systematically in Section 3.4.

3.4 Stage 1: Controlled Initial-Condition Placement

Now that the basics of feedback control have been reviewed, the remainder of this chapter explains how it is leveraged to construct basins of attraction experimentally. Stage 1 drives the system from a fixed rest position at the origin to the neighbourhood of the target cell $(x^*, \dot{x}^*)$, using law (21) to track a reference orbit that passes through the target state by construction. Once the actual state enters the target cell, the controller is switched off and Stage 2 begins.

3.4.1 Sinusoidal Reference Trajectory

A sinusoidal reference $x_r(t) = A\sin(\omega_r t + \phi_0)$ passing through $(x^*, \dot{x}^*)$ requires amplitude

$$A = \sqrt{x^{*2} + \frac{\dot{x}^{*2}}{\omega_r^2}}, \tag{22}$$

with initial phase $\phi_0 = \text{atan2}(x^*, \dot{x}^*/\omega_r)$. This formula allocates mechanical energy as in a simple harmonic oscillator, ignoring the cubic term βx^3, so the reference orbit is an exact ellipse in the phase plane regardless of the actual Duffing nonlinearity. The frequency ω_r is a free design parameter controlling the aspect ratio of the ellipse; all choices pass through $(x^*, \dot{x}^*)$ by construction, as illustrated in Figure 17.

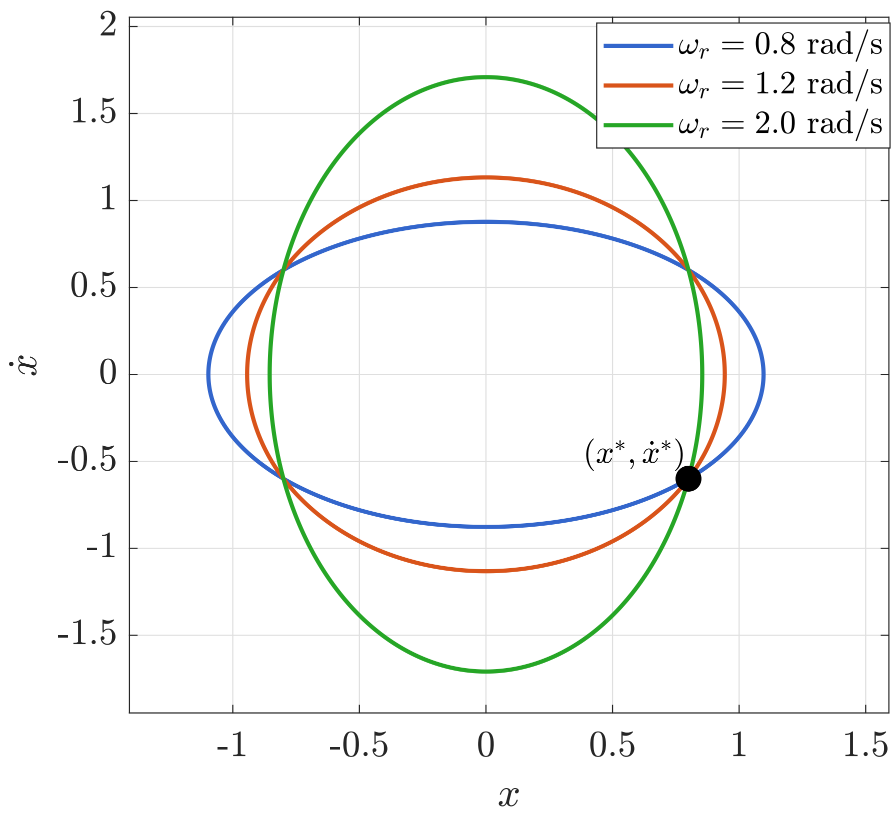

Figure 17: Sinusoidal reference orbits for three values of ω_r, all passing through the same target state (filled circle). The frequency ω_r controls the orbit aspect ratio without affecting the target.

The limitation of the sinusoidal reference is energetic inconsistency with the Duffing oscillator when the nonlinearity is significant: the reference lies on an energy level that differs from the true Duffing energy E^* of the target state. Intuitively, the Duffing oscillator naturally follows trajectories shaped by the elliptic functions introduced in Section 3.4.3; forcing it to follow a purely sinusoidal reference requires a sustained control effort throughout Stage 1. The magnitude of this mismatch is characterised by the elliptic modulus introduced below.

3.4.2 Validity Range of the Sinusoidal Reference

The energetic error is characterised by the elliptic modulus

$$m' = \frac{\varepsilon A^2}{2\big(1+\varepsilon A^2\big)}, \qquad \varepsilon = \frac{\beta}{\alpha}, \tag{23}$$

evaluated at the sinusoidal amplitude (22). The modulus m' quantifies the departure of the true Duffing orbit from a sinusoid: for $m' \approx 0$ the sinusoidal reference is adequate, while as m' approaches $\frac{1}{2}$ the approximation breaks down. The critical amplitude above which the approximation error exceeds a chosen threshold m'_0 is

$$A_{\text{crit}}(m'_0) = \sqrt{\frac{2m'_0}{\varepsilon(1-2m'_0)}}. \tag{24}$$

The threshold $m'_0 = 0.1$ is adopted throughout. For target cells satisfying $A < A_{\text{crit}}(0.1)$ the sinusoidal reference is used; for cells beyond this threshold the elliptic reference of Section 3.4.3 is required.

Figure 18 illustrates this architectural departure in the state space. When $m' \to 0$, the phase orbit traces a perfect ellipse corresponding to linear harmonic motion. As the modulus increases, the orbit evolves into a characteristic stadium-like shape. This state-space distortion highlights the structural error introduced if a pure sinusoidal reference is imposed on high-energy Duffing orbits.

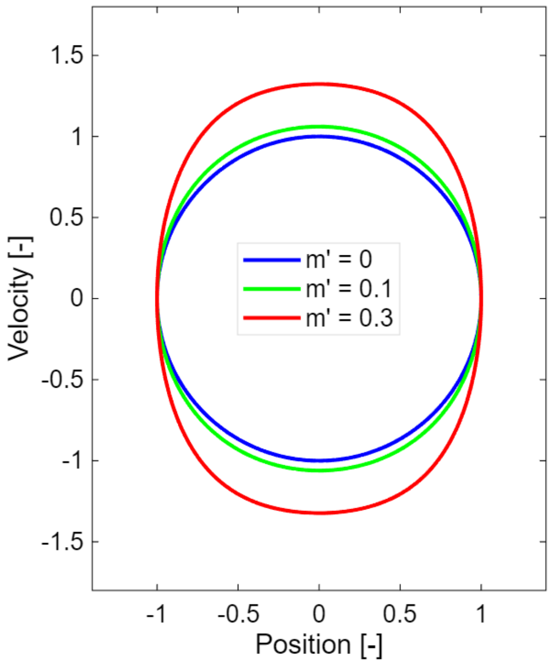

Figure 18: State-space phase orbits showing the transition from a linear harmonic trajectory to Duffing ones.

3.4.3 Non-Sinusoidal Reference: Jacobi Elliptic Functions

The energetic inconsistency is eliminated by replacing the sinusoidal reference with an exact solution of the conservative Duffing equation $\ddot{x} + \alpha x + \beta x^3 = 0$. Multiplying by $\dot{x}$ and integrating gives the first integral $E = \frac{1}{2}\dot{x}^2 + \frac{\alpha}{2}x^2 + \frac{\beta}{4}x^4$. Evaluating at the target state defines E^*; solving for the turning amplitude gives the corrected amplitude

$$A = \sqrt{\frac{-\alpha + \sqrt{\alpha^2 + 4\beta E^*}}{\beta}}, \tag{25}$$

which reduces to (22) as $\beta \to 0$. A reference built from this amplitude lies exactly on the energy level E^*, so no sustained energy input is required from the controller. The exact periodic solutions are Jacobi elliptic functions whose form depends on α and the position of E^* relative to E_s, as summarised in Table 1.

An inadequate reference can be problematic in strongly nonlinear regimes of motion. For the Duffing oscillator considered in this work, using the Jacobi elliptic functions as a reference was observed to result in much better tracking performance. Finding adequate references for more general systems is left for future work.

Table 1: Reference trajectories of the conservative Duffing oscillator. The phase t_0 is obtained by inverting the elliptic function at x^*; its sign is chosen to match $\mathrm{sgn}(\dot{x}^*)$. The cn frequency and modulus are $\Omega = \sqrt{\alpha + \beta A^2}$ and $m' = \beta A^2/(2(\alpha + \beta A^2))$; the dn frequency and modulus are $\Omega_{\mathrm{dn}} = A\sqrt{\beta/2}$ and $m'' = -2\alpha/(\beta A^2)$.

Regime	**Condition**	**Solution** $x_r(t)$	**Modulus**
$\alpha > 0$, all orbits		$A\,\mathrm{cn}(\Omega(t+t_0) \mid m')$	$m' \in (0, \frac{1}{2})$
$\alpha < 0$, cross-well	$E > E_s$	$A\,\mathrm{cn}(\Omega(t+t_0) \mid m')$	$m' \in (\frac{1}{2}, 1)$
$\alpha < 0$, separatrix	$E = E_s$	$A\,\mathrm{sech}(\sqrt{-\alpha}(t+t_0))$	$m' = 1$
$\alpha < 0$, in-well	$E < E_s$	$\pm A\,\mathrm{dn}(\Omega_{\mathrm{dn}}(t + t_0) \mid m'')$	$m'' \in (0, 1)$

The in-well reference is multiplied by $\mathrm{sgn}(x^*)$ to select the correct potential well. For in-well cells the sinusoidal reference is not merely inaccurate but structurally inadmissible: no sinusoid centred at the origin can lie on the energy level of a cell located within one well, regardless of amplitude or frequency. The dn reference is the only physically admissible option.

Figure 19 compares elliptic and sinusoidal orbits for six representative initial conditions.

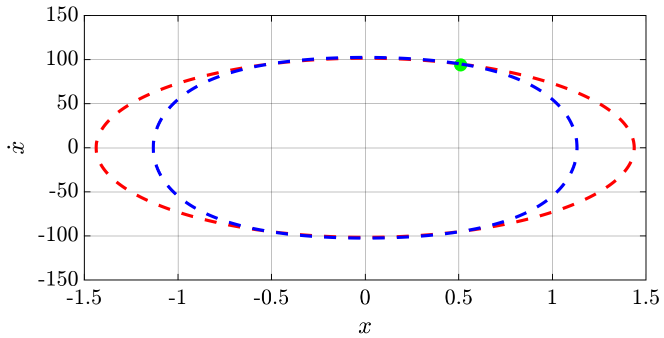

(a) IC-M1 (monostable): small mismatch.

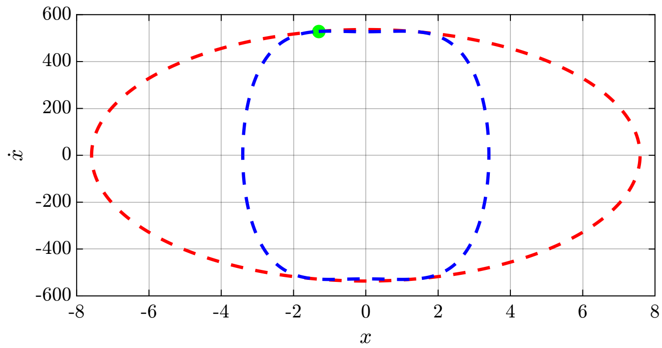

(b) IC-B1 (bistable, out-of-well): large mismatch.

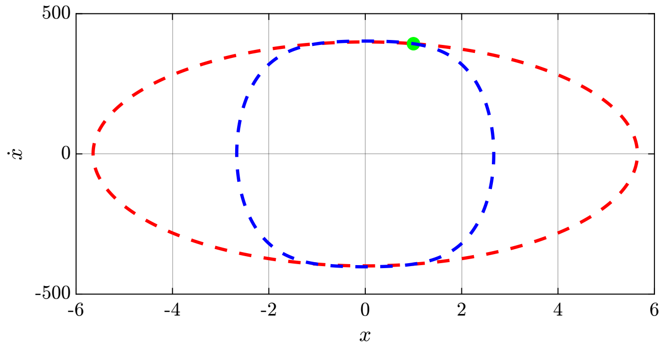

(c) IC-M2 (monostable): moderate mismatch.

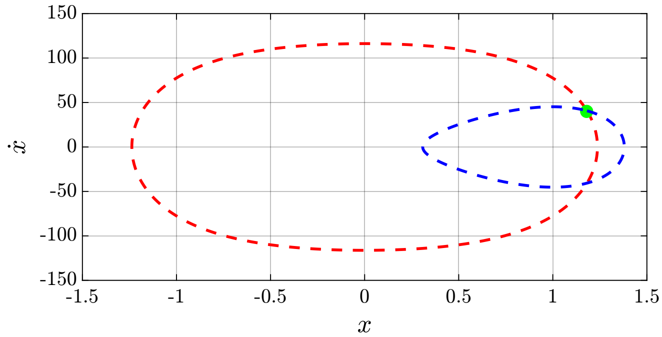

(d) IC-B2 (bistable, in-well): structural mismatch.

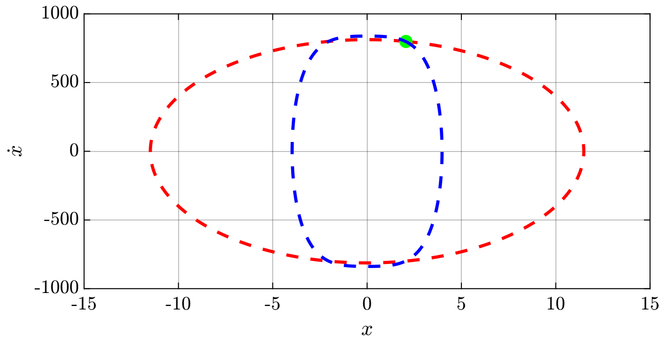

(e) IC-M3 (monostable): large mismatch at high amplitude.

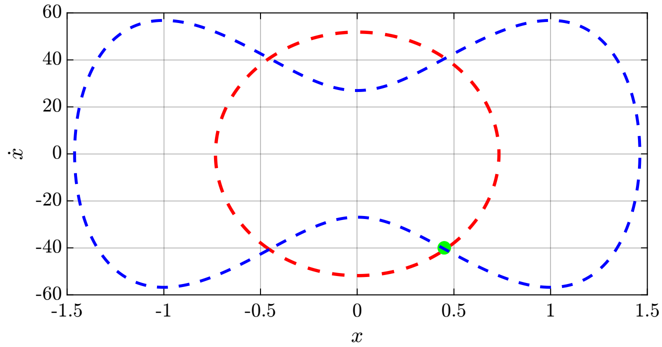

(f) IC-B3 (bistable, out-of-well): large mismatch.

Figure 19: Exact elliptic orbit (dashed blue) and sinusoidal approximation (dashed red) for six initial conditions. The mismatch grows with amplitude for monostable targets and is structural for bistable in-well cases.

3.4.4 Smooth Reference Engagement

When the controller is first activated, the system is at rest while the reference orbit oscillates at full amplitude. An abrupt engagement would produce an impulsive control force and risk actuator saturation. The reference is therefore multiplied by a raised-cosine ramp

$$G(t) = \tfrac{1}{2}\big[1 - \cos(\pi t/T_r)\big], \quad 0 \le t < T_r; \qquad G(t) = 1, \quad t \ge T_r, \tag{26}$$

which satisfies $G(0) = 0$, $G(T_r) = 1$, $\dot{G}(0) = \dot{G}(T_r) = 0$, so no velocity discontinuity is introduced at either end.

3.4.5 Gain Tuning and Scheduling

Substituting (21) into the Duffing equation and collecting error terms, and neglecting the nonlinear term relative to the proportional restoring force (a valid assumption when the closed-loop bandwidth ω_c satisfies

$\omega_c \gg \sqrt{\beta}A$), gives the closed-loop error dynamics

$$m\ddot{e} + (c + K_d)\dot{e} + (k + K_p)e = 0. \tag{27}$$

Comparing with $\ddot{e} + 2\zeta\omega_c\dot{e} + \omega_c^2 e = 0$ yields

$$K_p = m\omega_c^2 - k, \qquad K_d = 2\zeta\omega_c m - c. \tag{28}$$

The Butterworth criterion $\zeta = 1/\sqrt{2}$ is adopted, minimising the $\mathcal{H}_2$ norm of the error for a second-order system [5].

The bandwidth is scheduled cell-by-cell as

$$\omega_c(A) = \kappa\,\Omega(A), \qquad \kappa = 5, \tag{29}$$

where $\Omega(A)$ is the backbone frequency from the cn or dn solution, as shown in Figure 20. The factor $\kappa = 5$ places the controller bandwidth above the oscillation frequency, ensuring that the tracking error decays within a fraction of one cycle. Since A depends on $(x^*, \dot{x}^*)$ through (25), the gains vary smoothly across the grid and are precomputed for every cell before the start of the two-stage procedure.

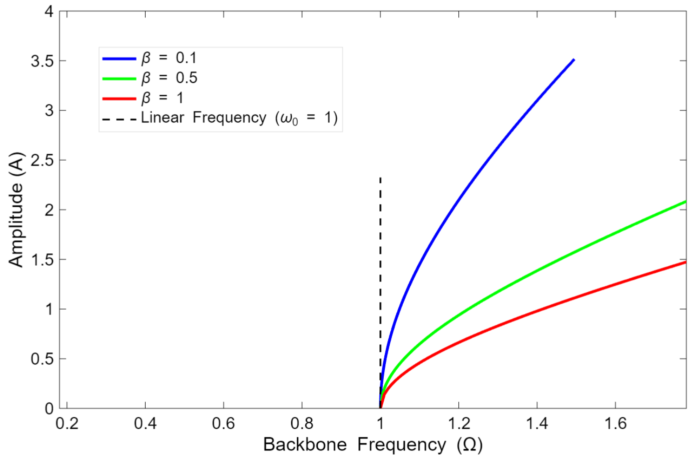

Figure 20: Backbone frequency $\Omega(A) = \sqrt{\alpha + \beta A^2}$ for several values of β/α. The gain scheduling ties the controller bandwidth to this curve, keeping the bandwidth-to-frequency ratio constant across all cells.

3.4.6 Cell-Entry Detection Criterion

The phase space is discretised into an $n \times n$ uniform grid covering $[-x_0, x_0] \times [-\omega_0 x_0, \omega_0 x_0]$, with cell half-widths $\Delta x = 2x_0/n$ and $\Delta\dot{x} = 2\omega_0 x_0/n$. Stage 1 declares the trajectory on target when the normalised distance to the cell centre satisfies

$$d(t) = \sqrt{\left(\frac{x - x^*}{\Delta x}\right)^2 + \left(\frac{\dot{x}/\omega_0 - \dot{x}^*/\omega_0}{\Delta\dot{x}/\omega_0}\right)^2} \le \tfrac{1}{2}. \tag{30}$$

The threshold $\frac{1}{2}$ corresponds to the inscribed circle of the target cell; the velocity axis is normalised by ω_0 to keep the two components dimensionally homogeneous. Figure 21 depicts the geometrical entities involved.

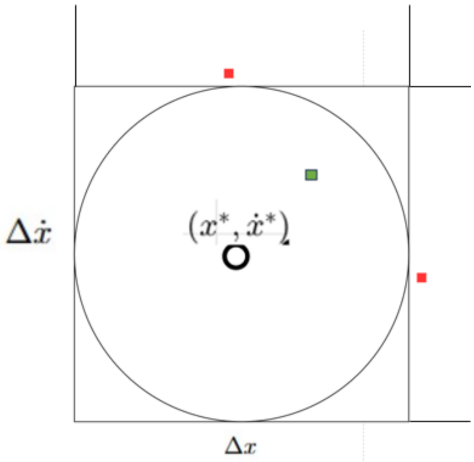

Figure 21: Cell-entry detection criterion. The shaded circle of radius $\frac{1}{2}$ (in normalised coordinates) is the inscribed circle of the target cell. Stage 1 triggers when the trajectory enters this circle.

3.4.7 Inter-Sample Interpolation

The criterion (30) is evaluated at discrete sampling times $t_k = k\Delta t$. Step-over occurs when the trajectory crosses the full cell diameter in one sampling step, i.e. when $|\dot{x}|\,\Delta t > \Delta x$. The remedy is inter-sample interpolation: at each step the minimum normalised distance from the line segment $[s_{k-1}, s_k]$ to the cell centre s^* is computed in closed form. If it falls below $\frac{1}{2}$, the trajectory is declared to have crossed the cell and the crossing time is estimated by linear projection:

$$t^* = t_{k-1} + \Delta t \cdot \frac{\langle s^* - s_{k-1},\; s_k - s_{k-1}\rangle}{\|s_k - s_{k-1}\|^2}, \tag{31}$$

where $s_k = (x(t_k)/\Delta x,\; \dot{x}(t_k)/(\omega_0 \Delta\dot{x}/\omega_0))$ is the normalised state at step k. This recovers sub-step accuracy at the cost of one additional stored state vector. The step-over problem and its resolution are illustrated in Figure 22.

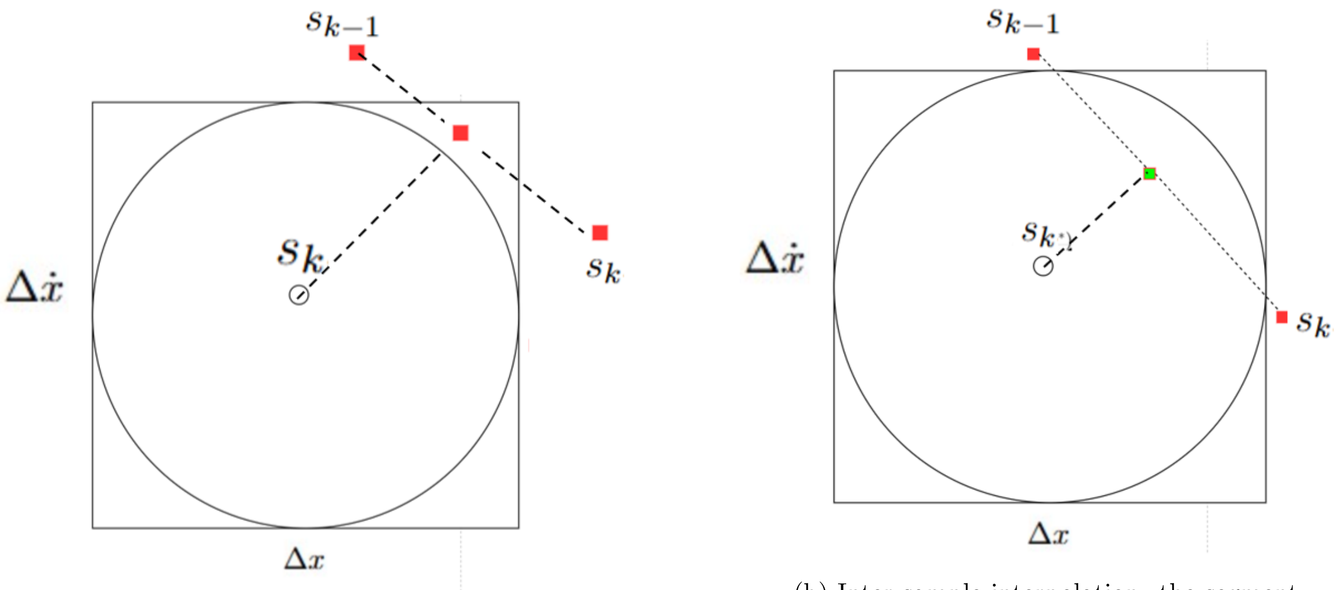

(a) Step-over failure: the trajectory crosses the cell between samples and the criterion never fires.

(b) Inter-sample interpolation: the segment minimum distance detects the crossing and the trigger fires correctly.

Figure 22: Illustration of the step-over problem and its resolution by inter-sample interpolation.

3.5 Stage 2: Free Evolution and Steady-State Identification

Once the state enters the target cell, the PD controller is switched off and the system evolves freely (autonomous case) or under the external forcing alone (forced case). The release state $(x_{\mathrm{rel}}, \dot{x}_{\mathrm{rel}})$ recorded at the trigger instant is the effective initial condition whose basin membership Stage 2 determines.

3.5.1 Stroboscopic Poincaré Sampling

Steady-state convergence is assessed through a stroboscopic Poincaré map: the system state is sampled at each integer multiple of a reference period T. For the autonomous case T is the natural period computed from the energy at the release state; for the forced case $T = T_f = 2\pi/\omega_f$, where T_f is the harmonic forcing period. The construction of the stroboscopic section is illustrated in Figure 23.

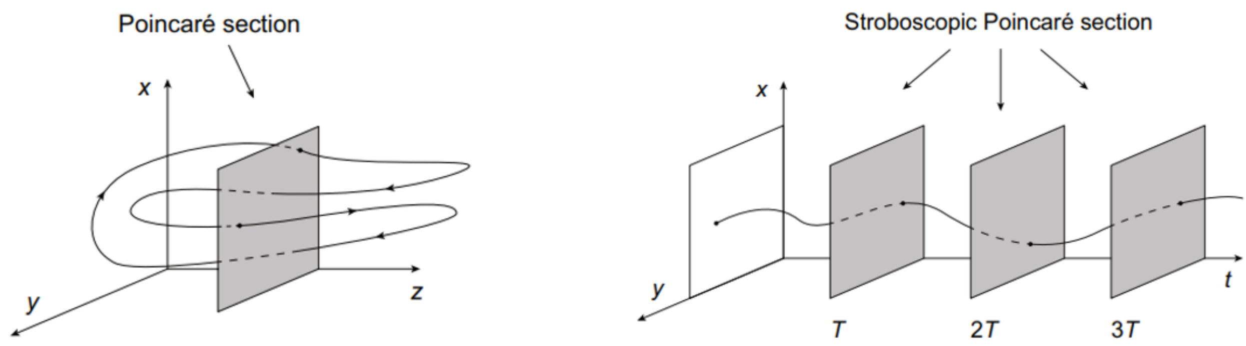

Figure 23: Stroboscopic Poincaré section construction [1].

3.5.2 Convergence Criterion

At each Poincaré instant the current state $(x, \dot{x})$ is compared to the n most recently stored samples using amplitude-normalised errors

$$\varepsilon_x^{(k)} = \frac{|x - x_k|}{A}, \qquad \varepsilon_{\dot{x}}^{(k)} = \frac{|\dot{x} - \dot{x}_k|}{\Omega A}, \tag{32}$$

where A and Ω are the amplitude and frequency at the release state. A match at depth k is declared when both normalised errors fall below τ_{pc}. The OR condition across all depths allows the criterion to accept period-p attractors for $p \in \mathbb{N}_0$: a period-p orbit returns to the same Poincaré point every p periods and satisfies the depth-p match condition even if it does not match at depth 1. Convergence is declared when n_{stable} consecutive Poincaré samples all satisfy the match condition:

$$\text{converged} \iff \exists\, k \in \{1, \dots, n\} : \varepsilon_x^{(k)} < \tau_{\mathrm{pc}} \text{ and } \varepsilon_{\dot{x}}^{(k)} < \tau_{\mathrm{pc}} \quad \text{for } n_{\mathrm{stable}} \text{ consecutive periods.} \tag{33}$$

Figure 24 illustrates convergence detection for a period-1 and a period-3 attractor.

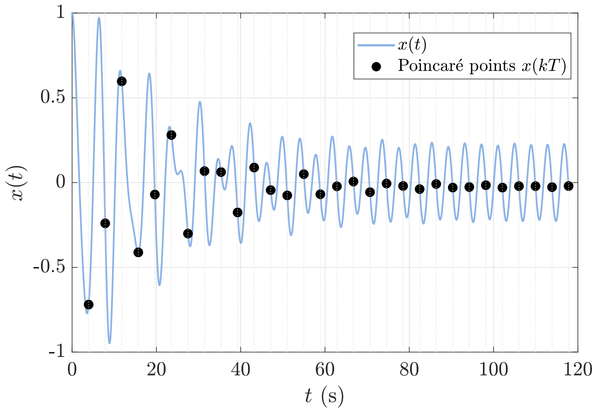

(a) Period-1 time response with Poincaré samples (black dots) converging to a single value.

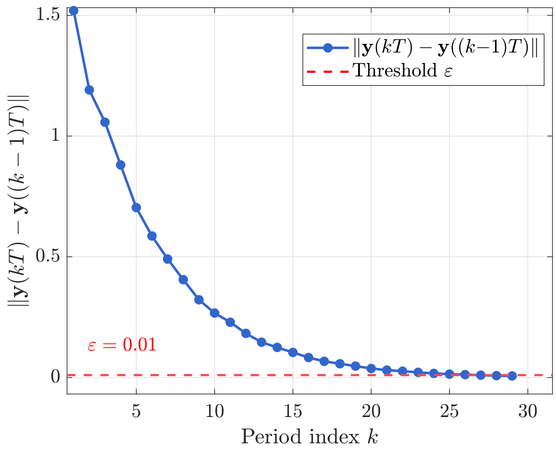

(b) Period-1 depth-1 error decaying below τ_{pc}.

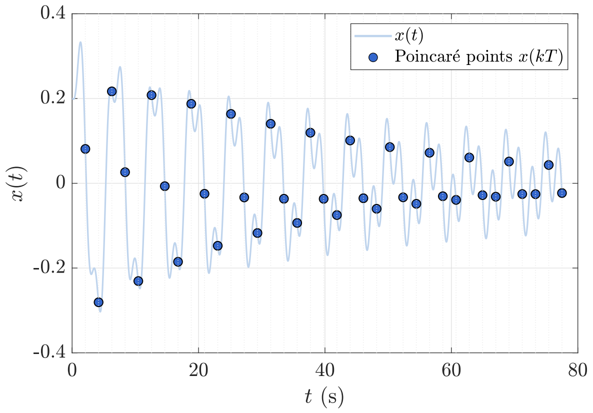

(c) Period-3 time response with Poincaré samples cycling through three distinct values.

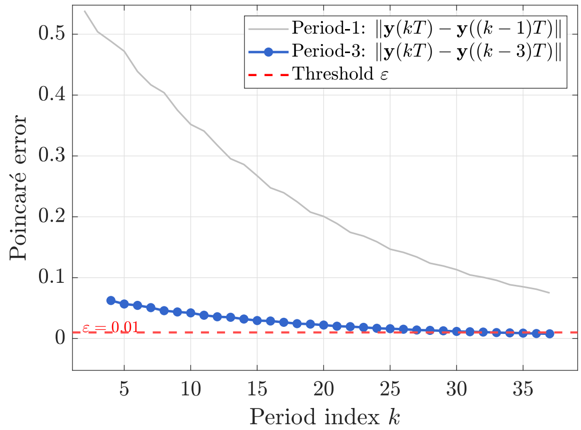

(d) Depth-1 error (grey) fails the threshold; depth-3 error (blue) clears it, correctly identifying a period-3 attractor.

Figure 24: Convergence detection using equation (33) for a period-1 attractor (top row, $\delta = 0.1$, $\beta = 0.3$, $\gamma = 0.35$, $\omega = 1.6$) and a period-3 attractor (bottom row, $\delta = 0.05$, $\beta = 1$, $\gamma = 0.65$, $\omega = 3$).

The sensitivity of the reconstructed basin to the pair $(\tau_{\text{pc}}, n_{\text{stable}})$ is studied in Section 4.5.2; the validated values $\tau_{\text{pc}} = 0.01$ and $n_{\text{stable}} = 20$ are adopted for all main results.

3.5.3 Attractor Label Assignment

For the autonomous system the attractor label is the stroboscopic position $x^{(\infty)}$ at convergence, identifying the equilibrium or periodic orbit reached. For the forced system the label is the peak amplitude

$$\rho = \max_{t \in [t_k, t_{k+1}]} |x(t)| \tag{34}$$

over the last confirmed period, which distinguishes the coexisting attractors and is the natural observable in a harmonic excitation experiment. Labels across all cells are post-processed by a nearest-neighbour clustering algorithm that groups cells within a fixed tolerance in attractor value, producing a small discrete set of centroid labels consistent across the full grid.

3.6 Basin Reconstruction Strategy

The procedure outlined thus far allows for the experimental determination of the attractor associated with a given initial condition. The missing ingredient to construct basins of attraction is to distribute initial conditions over the state space. This final section discusses this aspect.

3.6.1 Spiral Sampling Order

Target states are visited along a discrete Archimedean (Fermat) spiral

$$x_k = \rho_x \sqrt{\theta_k} \cos\theta_k, \qquad \dot{x}_k = \rho_v \sqrt{\theta_k} \sin\theta_k, \tag{35}$$

where ρ_x and ρ_v are scaling factors set by the domain boundaries. The square-root radial term distributes sample points uniformly in area, avoiding over-density near the origin. Each analytical point is snapped to the nearest grid node and duplicate indices are removed. The spiral order allows early qualitative insight into the basin shape. The resulting sampling pattern is shown in Figure 25.

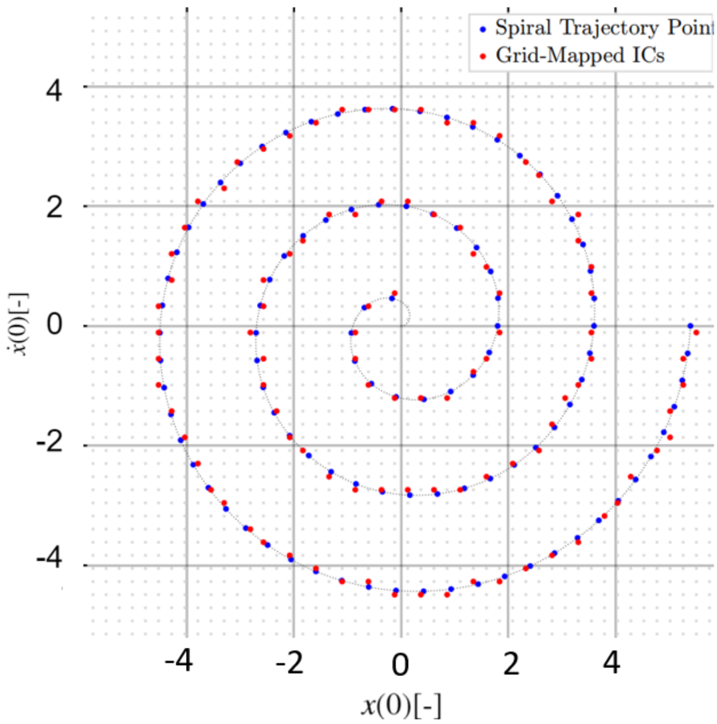

Figure 25: Spiral sampling pattern on a 100 × 100 grid. Cells are visited outward from the centre; the colour encodes the visit order from dark (first) to light (last).

Once the set of initial conditions is determined and the control gains are precomputed, the two-stage procedure is repeated for every initial condition, allowing for the experimental construction of basins of attraction. The proposed methodology is validated numerically in Chapter 4 and applied experimentally in Chapter 5.

3.6.2 Trajectory-Based Label Backfilling

The two-stage procedure assigns one attractor label per grid cell, but integrating every cell independently introduces significant redundancy for autonomous or periodically-forced systems: if a trajectory launched from one cell passes through the stroboscopic states of another, both cells necessarily share the same asymptotic behaviour. To exploit this, an additional feature inspired from cell-mapping algorithms [6] can be implemented. The algorithm records the stroboscopic Poincaré states visited during Step 2. Once convergence to an attractor is declared, these recorded states are used to retroactively label every grid cell they intersect, eliminating the need to launch subsequent two-stage procedures for those cells.

The efficiency of this backfill strategy depends on the cell-visiting order. Trajectories launched from high-energy outer cells traverse large regions of the phase plane before settling, and their stroboscopic paths cross many inner cells. Integrating outer cells first therefore maximises the number of inner cells that can be labelled via backfill. This motivates the concentric square sampling order introduced in Section 4.7: the outermost cells are integrated first, allowing their long-horizon trajectories to propagate labels inward and reduce the total computational overhead.

4 Numerical Validation

The two-stage controlled basin-of-attraction reconstruction algorithm is validated numerically across four Duffing configurations of increasing dynamical complexity. Each case is chosen to isolate and stress-test a specific aspect of the reconstruction pipeline, so that the progression builds a comprehensive picture of the method's capabilities and limits before experimental deployment.

The four configurations are summarised in Table 2. Case 1 is a weakly nonlinear forced system ($\varepsilon = 0.052$) operating within a narrow bistable band. Its primary role is to establish the sinusoidal reference as a valid and efficient choice for low-amplitude initial conditions, and to identify the amplitude threshold beyond which it breaks down. Case 2 raises the nonlinearity ratio to $\varepsilon = 1$, widening the bistable band considerably and pushing the system well outside the sinusoidal validity domain. This case demonstrates that the elliptic reference faithfully reconstructs the basin where the sinusoidal one fails. Case 3 removes the forcing entirely and reverses the sign of the linear stiffness, producing an autonomous double-well system whose two attractors are fixed points rather than periodic orbits. Because the sinusoidal reference is structurally inapplicable here, this case specifically validates the dn/cn elliptic switching strategy that handles in-well and cross-well dynamics within a single unified framework. Case 4 combines very low damping with a large forcing amplitude, giving rise simultaneously to fundamental bistability, $1:2$ and $1:3$ subharmonic resonances, and an isolated response branch. It represents the most demanding test of the algorithm, confirming that the controlled reconstruction correctly identifies each coexisting attractor even when the basin geometry is highly intricate.

Table 2: Four Duffing configurations for numerical validation. The nonlinearity ratio $\varepsilon = \beta/|\alpha|$ governs the validity of the sinusoidal reference.

Case	α	β	δ	γ	ν	**Mode**
1: weakly NL forced	1	0.052	0.016	0.039	1.046	Forced
2: strongly NL forced	1	1	0.1	0.35	1.6	Forced
3: unforced bistable	−1	1	0.1	0	—	Autonomous
4: highly forced	1	1	0.05	3	—	Forced

Before discussing individual cases, two practical aspects of the algorithm must be addressed. In Stage 1, the controlled trajectory may not always reach the desired neighbourhood of the target initial condition within a finite time, for instance when the reference tracking becomes problematic at very high amplitudes. A timeout parameter $t_{\max,1}$ is therefore introduced: if Stage 1 has not triggered by $t_{\max,1}$, the cell is marked as not converged (NC) and the algorithm moves to the next initial condition without releasing the system. In Stage 2, excessively long transients can occur in exceptional cases, for instance near a fold bifurcation where critical slowing down is present. A maximum number of stroboscopic periods $N_{\max}$ is set; if this limit is reached before the convergence criterion of equation (33) is satisfied, the cell is similarly marked as NC. In practice, this situation arises when an initial condition lies in the vicinity of the basin boundary, where the trajectory may linger near the unstable branch for an extended duration. NC cells are displayed in black in all basin maps, and in practice this situation generally arises over thin regions near the basin boundary. A critical analysis of results in those regions is therefore required.

Algorithm parameters common to all four cases are listed in Table 3; their selection is justified in Section 4.5.2. For each case, the presentation follows a fixed structure: the frequency-response curve (FRC) and elliptic modulus analysis establish the operating context, the exact basin obtained by direct integration provides the ground truth, and the controlled basin together with its point-wise difference map quantifies the reconstruction accuracy. Beyond the individual cases, Section 4.4 traces how the basin evolves continuously across the bistable band for Cases 1 and 2, illustrating attractor dominance transitions as a function of excitation frequency. Section 4.5 then quantifies the sensitivity of the reconstructed basin to physical parameter uncertainty and to the choice of algorithmic thresholds. Section 4.6 concludes the case studies with the multi-branch analysis of Case 4. Finally, Section 4.7 demonstrates that the trajectory-based backfill acceleration reduces active integration to 30% of the grid across all four configurations while preserving full topological fidelity. Simulations were performed across 32 CPU cores on a cluster of the CÉCI consortium [3]. [1]

As the spiral sampling order introduced in Section 3.6.1 is primarily motivated by the experimental context, where qualitative insight into the basin structure is desirable early in the sweep, the numerical validation uses a standard rectangular grid ordering. The computational resources available on the cluster made it unnecessary to exploit the spiral order for the simulations.

Table 3: Algorithm parameters common to all four cases. τ_{pc} and n_{stable} are fixed at the values validated in Section 4.5.2.

Parameter	**Symbol**	**Value**	**Role**
Grid density (per axis)	n	100	Phase-space resolution
Stage 1 timeout	$t_{\max,1}$	20 s	Maximum approach time
Stage 1 step	Δt_1	10^{-4} s	Fixed ODE step
Poincaré tolerance	τ_{pc}	0.01	Convergence test
Stable-count	n_{stable}	20	Confidence filter
Period limit	$N_{\max}$	200	Stage 2 timeout

4.1 Case 1: Weakly Nonlinear Forced Bistable System

4.1.1 Frequency-Response Curve

Figure 26 shows the harmonic balance FRC obtained using the continuation framework of [2] for Case 1. The system exhibits classical hardening behaviour, characterised by a rightward tilt of the resonance peak. This geometric distortion gives rise to a bistable region bounded by two fold (saddle-node) bifurcations. Within this coexistence zone, a high-amplitude resonant branch and a low-amplitude non-resonant branch coexist alongside an intermediate unstable branch.

The nominal operating frequency is set to $\nu_{\mathrm{op}} = 1.046$, placing it within the lower portion of the bistable band. This regime is deliberately chosen because the upper stable branch amplitude is still relatively low, allowing bistability to be observed even for low-amplitude initial conditions.

Operating in this region introduces a trade-off regarding transient dynamics. Due to its proximity to the lower fold bifurcation ($\nu \approx 1.03$), the system operates near a phase-space bottleneck, a phenomenon known as critical slowing down. A trajectory may oscillate in the vicinity of the unstable branch for an extended duration before slowly settling to steady state. In practice, this trade-off thus occurs between an accurate basin reconstruction and the computational complexity of the method, which translates to test duration in an experiment.

[1] Computational resources have been provided by the Consortium des Équipements de Calcul Intensif (CÉCI), funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under Grant No. 2.5020.11 and by the Walloon Region.

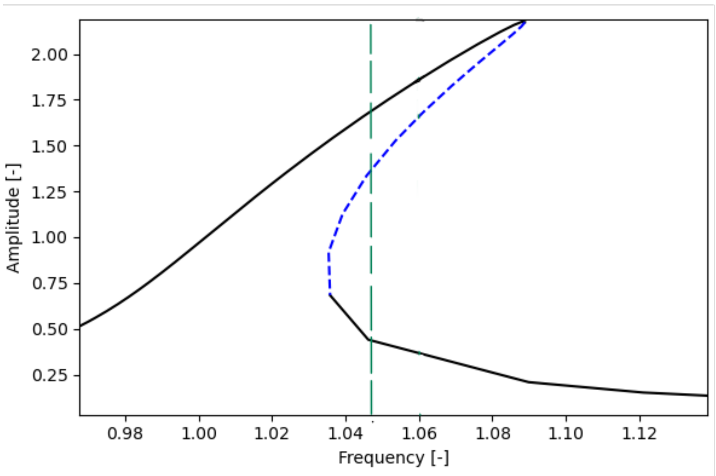

Figure 26: HB1 FRC for Case 1 ($\alpha = 1$, $\beta = 0.052$, $\delta = 0.016$, $\gamma = 0.039$). Solid: stable branches. Dashed: unstable branch. Green dashed vertical line: operating frequency $\nu_{\text{op}} = 1.046$.

4.1.2 Elliptic Modulus Analysis and Grid Selection

To reconstruct the basin using a sinusoidal reference, the maximum permissible displacement amplitude must first be established. The critical amplitude A_{crit} is calculated via (24) with $m'_0 = 0.1$. Figure 27 shows this threshold for Case 1, demonstrating that the sinusoidal approximation remains accurate ($m' < 0.1$) over a relatively large displacement span. Mapping these contours onto the phase plane in Figure 28 confirms that a grid spanning $[-1, 1] \times [-1, 1]$ can be selected with confidence: the attractor positions identified from Figure 26 lie well inside the $m' = 0.1$ contour.

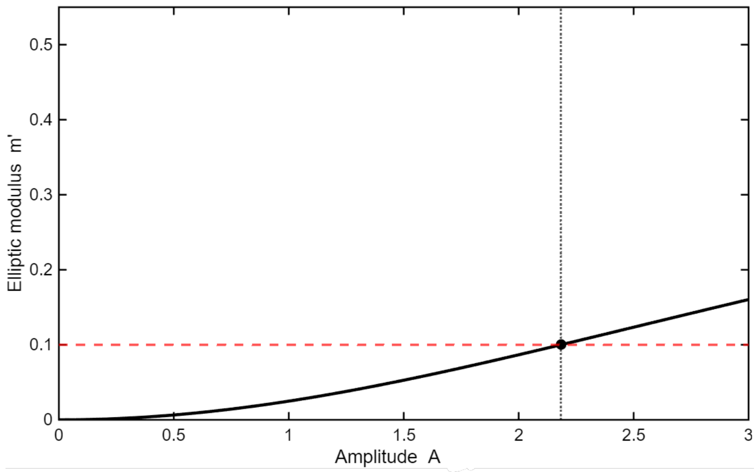

Figure 27: Elliptic modulus m' versus sinusoidal amplitude A_{sin} for Case 1 ($\varepsilon = 0.052$). The red dashed line marks $m' = 0.1$; A_{crit} is the intersection point.

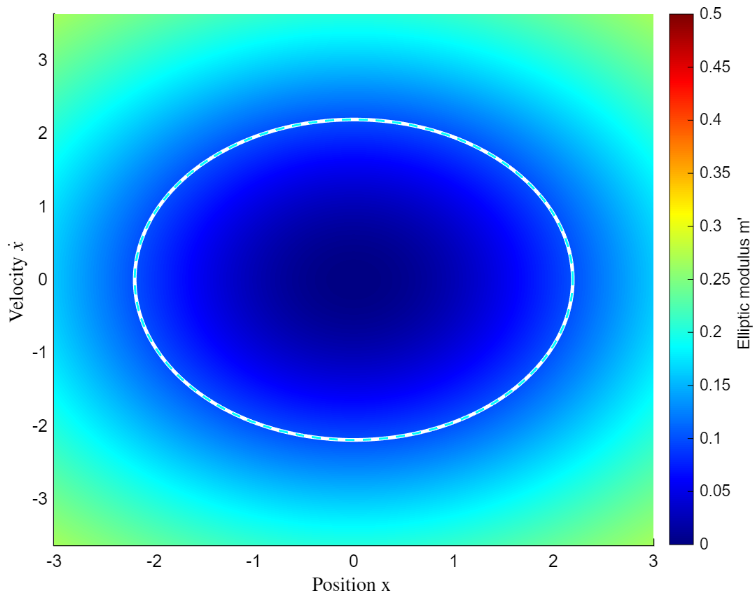

Figure 28: Elliptic modulus field $m'(x^*, \dot{x}^*)$ over the phase-space grid for Case 1. The white contour marks $m' = 0.1$.

4.1.3 Reference Basin of Attraction: Direct Integration

The reference basin, obtained by running direct Stage 2 integration from every grid cell without any active control loop, is shown in Figure 29. Two primary regions are visible, corresponding to the stable low- and high-amplitude regimes. Their steady-state amplitudes are entirely consistent with the respective branch responses computed from the FRC at $\nu_{\mathrm{op}} = 1.046$.

However, an anomalous feature emerges along the boundary separating the two primary domains. A subset of initial conditions located close to this boundary appears to converge to a distinct third attractor whose amplitude coincides with that of the intermediate unstable branch. This is a manifestation of the critical slowing down phenomenon discussed in Section 4.1.1.

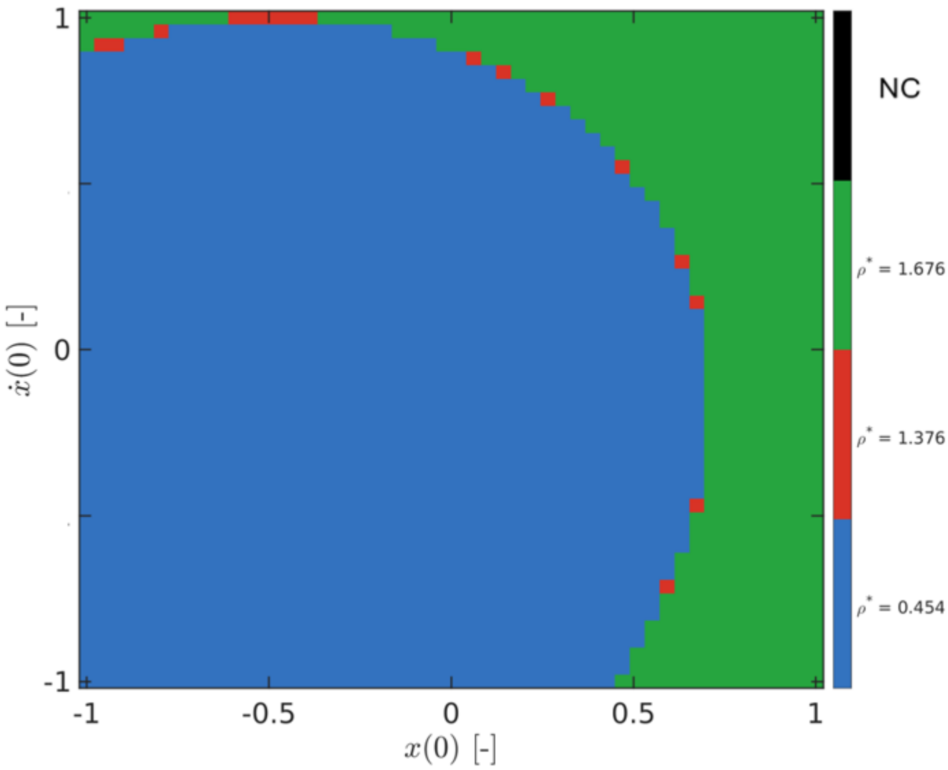

Figure 29: Exact basin for Case 1 via direct Stage 2 integration (no controller). Two primary stable domains are visible, as well as an anomalous third region along the boundary (NC: initial conditions that did not converge within the period limit).

To isolate the cause of this behaviour, the convergence speed map in Figure 30 tracks the number of periods required to declare steady state across the grid. Initial conditions near the boundary generally require more periods to settle. Notably, the cells that apparently converge to the unstable branch record a low convergence period count, lower than the surrounding boundary cells and than most of the phase plane. This confirms that the slow transients characteristic of critical slowing down trigger the convergence criterion prematurely: because the trajectory amplitude changes only minimally over consecutive periods within the phase-space bottleneck, the algorithm interprets this reduced variance as steady state before the trajectory has truly settled.

In an experiment, adopting stricter tolerances to mitigate this effect may be unfeasible due to measurement noise. Anticipating this situation, the tolerances used here are kept at reasonable levels. A critical analysis of the results near the basin boundary is thus required, but is eased by the fact that these artificial regions generally appear over thin strips near the boundary.

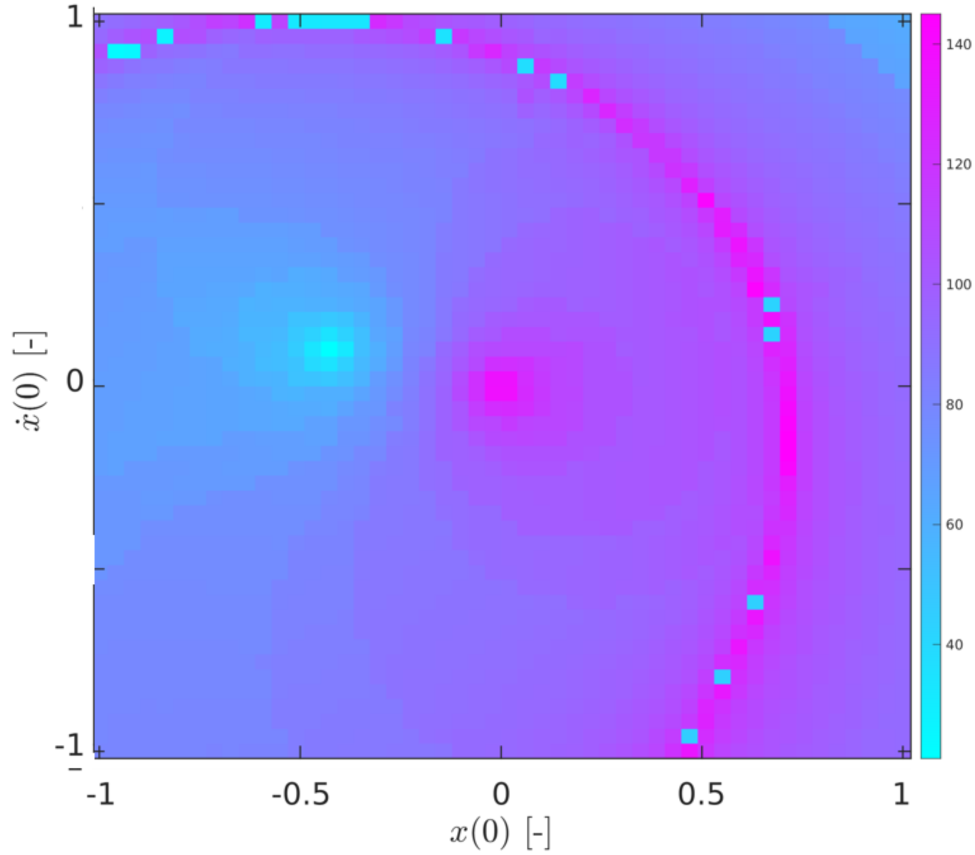

Figure 30: Convergence speed map for Case 1, showing the number of integration periods required to meet the steady-state criterion across the phase plane.

4.1.4 Controlled Basin of Attraction: Sinusoidal Reference

Gain scheduling. The PD gains are scheduled following (29) with $\kappa = 5$ and $\zeta = 1/\sqrt{2}$. As shown in Figure 31, both gains increase outward from the origin, keeping the bandwidth-to-frequency ratio constant across the grid.

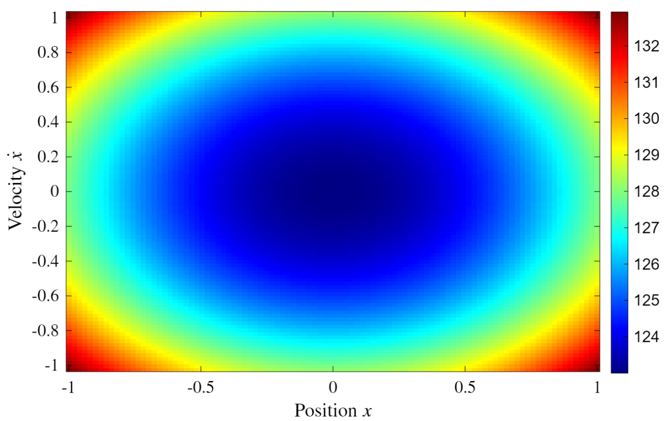

(a) Proportional gain K_p.

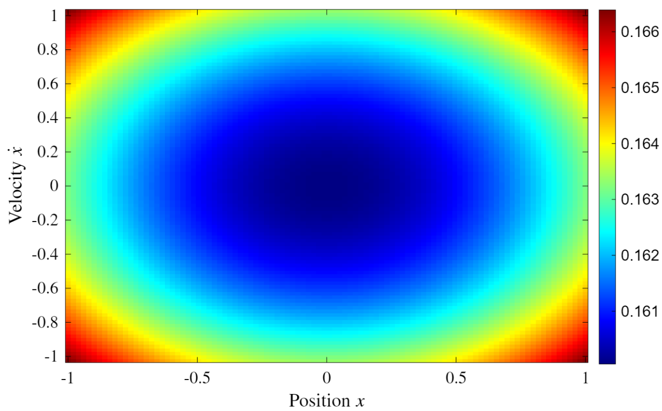

(b) Derivative gain K_d.

Figure 31: Scheduled PD gains for Case 1 (sinusoidal reference, $\kappa = 5$, $\zeta = 1/\sqrt{2}$).

Results on the inner grid. Figure 32(a) shows that the controller successfully drives the system to each target initial condition within the inner grid, with no NC cells. The controlled basin and the exact reference map are highly similar, with both primary regions and the separatrix clearly visible.

To evaluate accuracy quantitatively, a cell-by-cell difference map is shown in Figure 32(b). The discrepancies are strictly localised along the boundary between the two basins: small release errors at the end of

Stage 1 can alter the trajectory near the separatrix, causing it to converge to the other attractor. A residual subset of initial conditions still converges to the artificial third region, which is a remaining consequence of the critical slowing down discussed above.

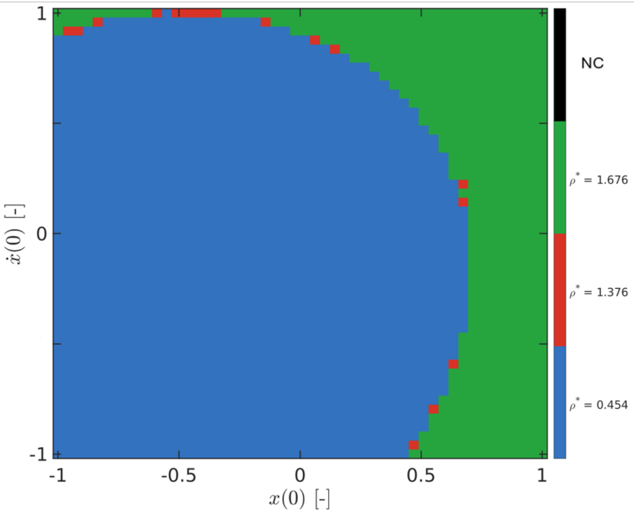

(a) Controlled basin (sinusoidal reference). Same colour convention as Figure 29.

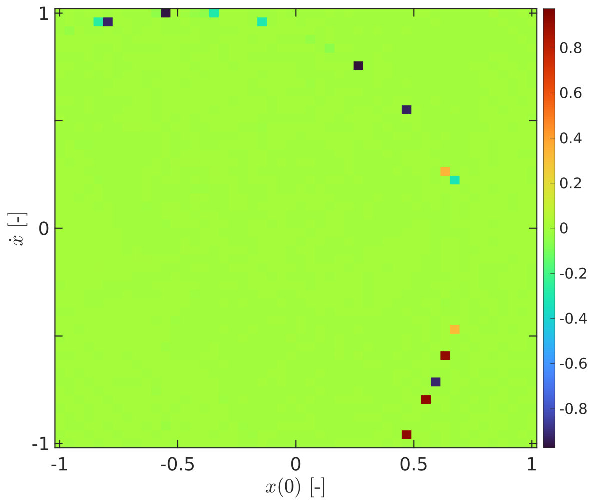

(b) Difference map $\Delta\rho$. Non-zero entries indicate misclassified cells.

Figure 32: Case 1 controlled basin (sinusoidal reference, inner grid) and point-wise comparison with the exact basin of Figure 29.

Breakdown on the extended grid. To evaluate the limits of the sinusoidal reference, the domain is extended while maintaining constant grid density. Figure 33(a) presents the exact basin on the extended grid; Figure 33(b) shows the controlled result.

While inner, low-amplitude initial conditions are correctly reached, a significant portion of the cells associated with high initial velocities result in Stage 1 NC (shown in black). Initial conditions with large displacements but low velocities remain reachable. This behaviour is consistent with the phase-orbit analysis of Figure 18: as the elliptic modulus increases, the phase-space trajectory adopts a stadium shape that diverges most strongly from a sinusoid at high velocity levels, while remaining relatively close near the zero-velocity axis. Because the sinusoidal reference does not accurately capture this distorted geometry, Stage 1 cannot guide the system to high-velocity cells within the timeout, resulting in NC classifications.

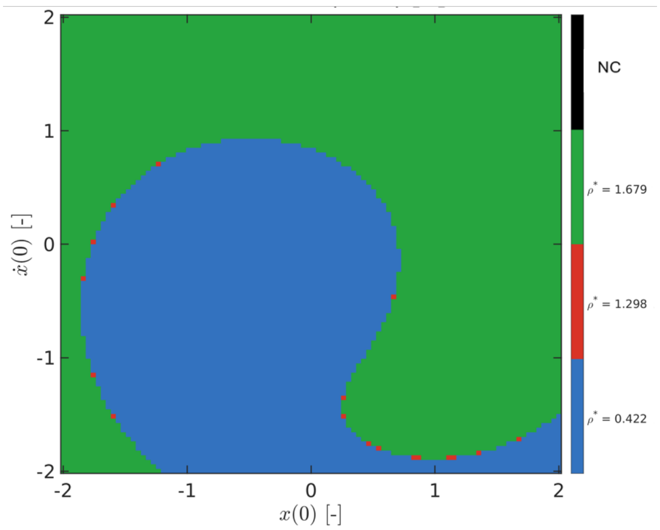

(a) Exact basin on the extended grid.

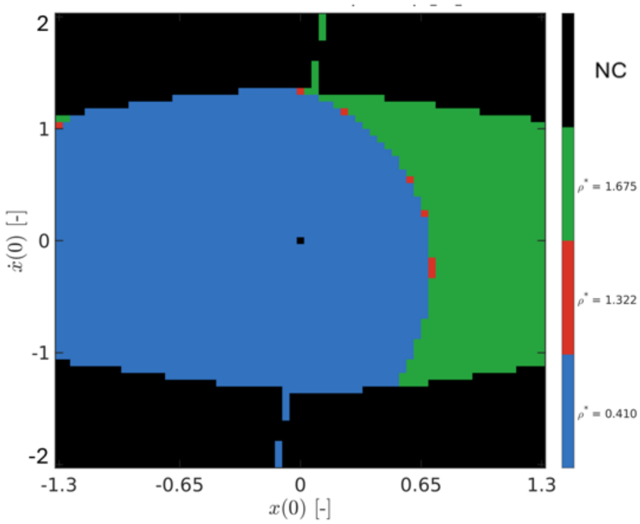

(b) Controlled basin (sinusoidal reference, extended grid). NC: cells where Stage 1 did not converge within $t_{\max,1}$.

Figure 33: Effect of extending the grid beyond A_{crit} for Case 1. The outer high-velocity region shows Stage 1 non-convergence with the sinusoidal reference.

4.1.5 Controlled Basin of Attraction: Elliptic Reference

Replacing the sinusoidal reference with the elliptic one resolves the outer-cell failures completely. As shown in Figure 34, the controlled basin now matches the exact basin throughout the extended grid, with residual differences confined to a single-cell-width strip along the basin boundary.

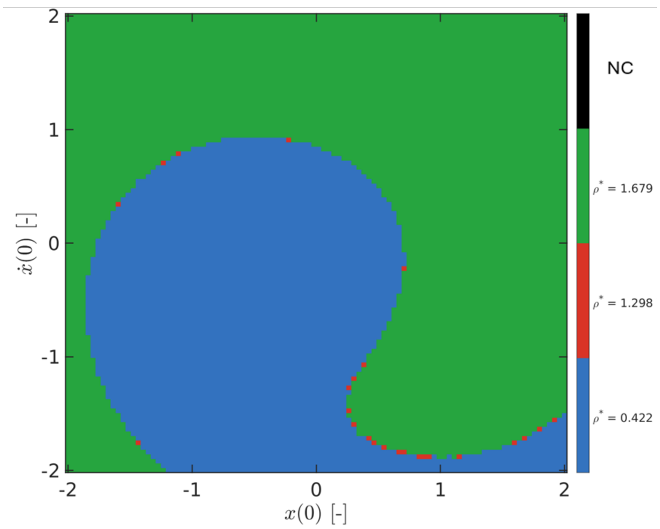

(a) Controlled basin (elliptic reference, extended grid). The NC region of Figure 33 is fully resolved.

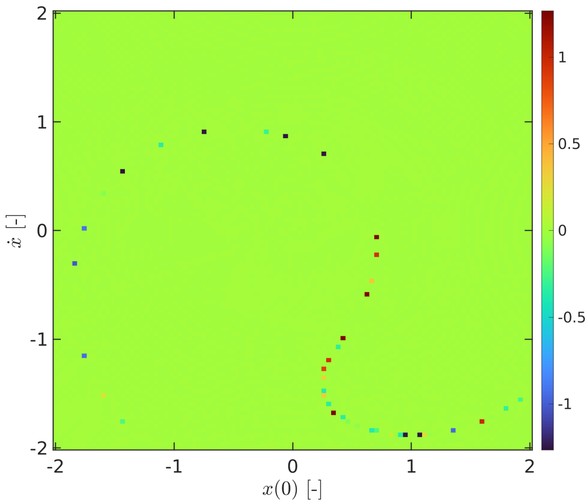

(b) Difference map $\Delta\rho$. Residual discrepancies are confined to a single-cell-width strip along the basin boundary.

Figure 34: Case 1 controlled basin with elliptic reference on the extended grid and comparison with the exact basin.

These results establish two benchmark conclusions. First, the sinusoidal reference remains adequate within the low-modulus region, in particular where initial velocities are small. However, high-velocity initial conditions accentuate the geometric discrepancy between the system's natural dynamics and a pure sinusoid, leading to Stage 1 tracking failures. Second, the elliptic reference corrects these localised failures outside A_{crit} by accurately capturing the energy level of every target cell, making it the reference of choice for

Cases 2 and beyond.

4.2 Case 2: Strongly Nonlinear Forced Bistable System

4.2.1 Frequency-Response Curve

Figure 35 shows the first-order harmonic balance FRC for Case 2. Due to the increased nonlinearity compared to Case 1, the resonance peak exhibits a more pronounced rightward tilt. This hardening behaviour widens the bistable region, which extends from approximately $\nu = 1.35$ to $\nu = 1.90$. The nominal operating frequency is fixed at $\nu_{\mathrm{op}} = 1.60$, placing the system in the middle of the bistable band.

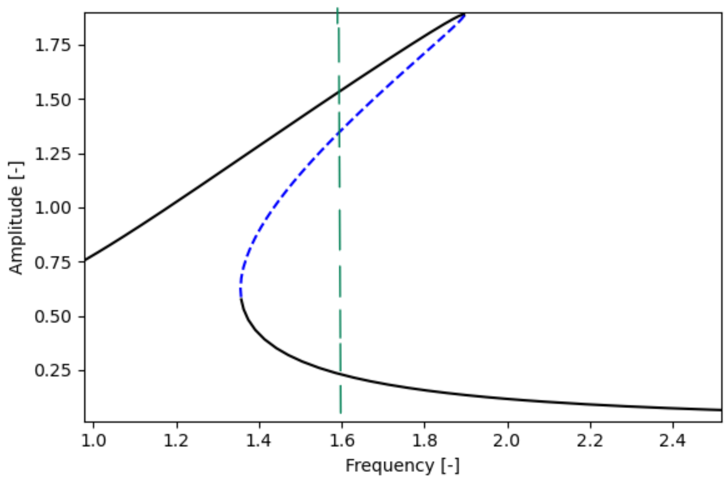

Figure 35: FRC for Case 2 ($\alpha = 1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0.35$). Same colour convention as Figure 26. The wider bistable band and larger amplitude difference between the two attractors reflect the larger nonlinearity ratio $\varepsilon = 1$.

4.2.2 Elliptic Modulus Analysis

Figure 36 overlays the modulus curves for Cases 1 and 2. The comparison shows that the critical amplitude threshold is substantially lower in Case 2, which stems from the larger β/α ratio. The steady-state trajectories therefore deviate significantly from a sinusoidal profile even for low-amplitude initial conditions, ruling out the sinusoidal reference across the full grid.

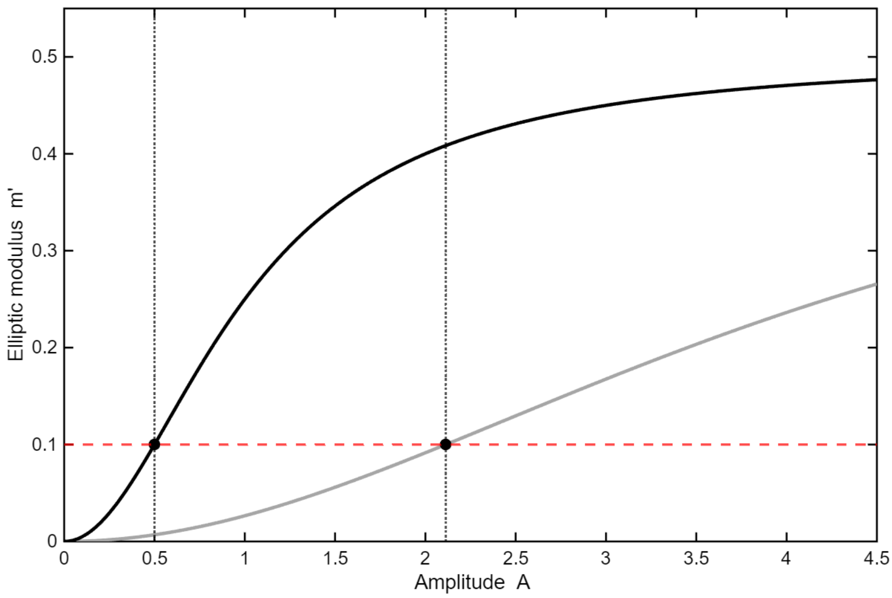

Figure 36: Elliptic modulus $m'(A_{\text{sin}})$ for Case 2 (black, $\varepsilon = 1$) and Case 1 (grey, $\varepsilon = 0.052$). Red dashed line: threshold $m' = 0.1$.

4.2.3 Reference Basin and Controlled Basin: Elliptic Reference

Gain scheduling. The gains for Case 2, shown in Figure 37, are higher than those in Case 1. This increase is a direct result of the hardening nonlinearity: the backbone frequency $\Omega(A)$ reaches higher values at large amplitudes, and the bandwidth scheduling $\omega_c = \kappa\,\Omega(A)$ scales the gains accordingly to maintain a consistent bandwidth-to-frequency ratio.

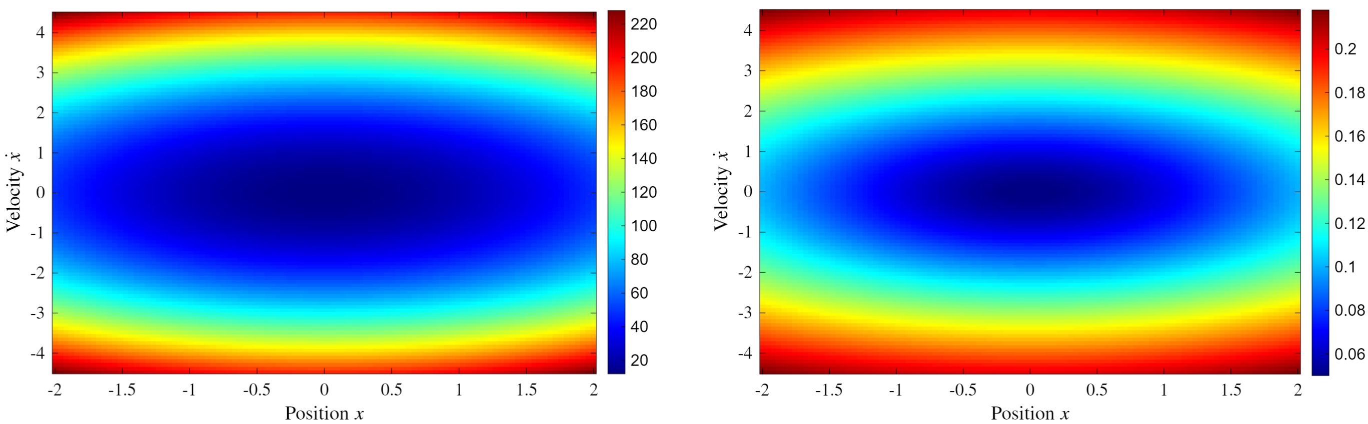

(a) Proportional gain K_p.

(b) Derivative gain K_d.

Figure 37: Scheduled PD gains for Case 2 (elliptic reference, $\kappa = 5$, $\zeta = 1/\sqrt{2}$). Gains are higher than in Case 1 (Figure 31) due to the larger backbone frequency at large amplitudes.

The four panels of Figure 38 present the exact basin, the controlled basin, the difference map, and the Stage 2 convergence speed. Because the operating frequency is located in the middle of the bistable band, neither attractor dominates the phase plane, resulting in an approximate 50/50 area proportion between the two basins.

The visual discrepancy between the controlled and exact basins is minor. The difference map reveals that mismatches are restricted to a narrow, single-cell-wide strip along the basin boundary, where small release errors can cause a trajectory to switch from one attractor to the other. The convergence speed map confirms that initial conditions close to the separatrix require more periods to settle than those within the basin interiors. Importantly, the convergence time remains lower than in Case 1, because the operating point in Case 2 is located further from both fold bifurcations and the effects of critical slowing down are therefore less pronounced.

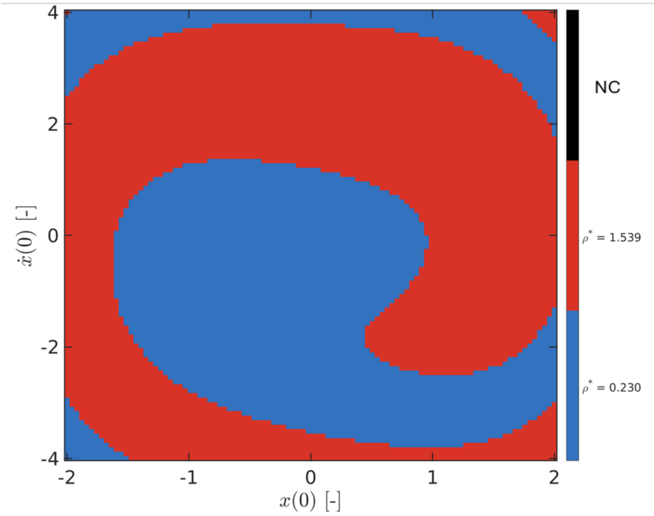

(a) Exact basin (direct integration, no controller).

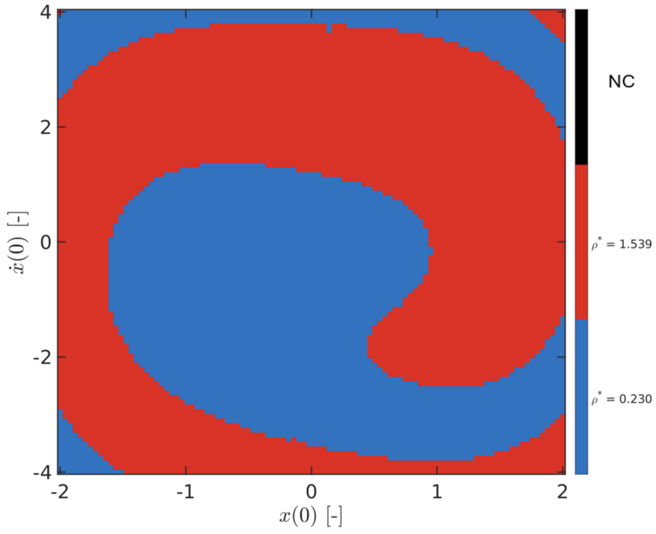

(b) Controlled basin (elliptic reference).

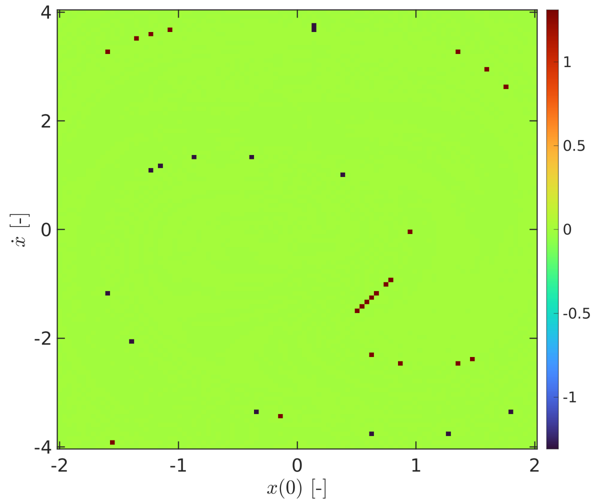

(c) Difference map $\Delta\rho$. Non-zero entries confined to the basin boundary strip.

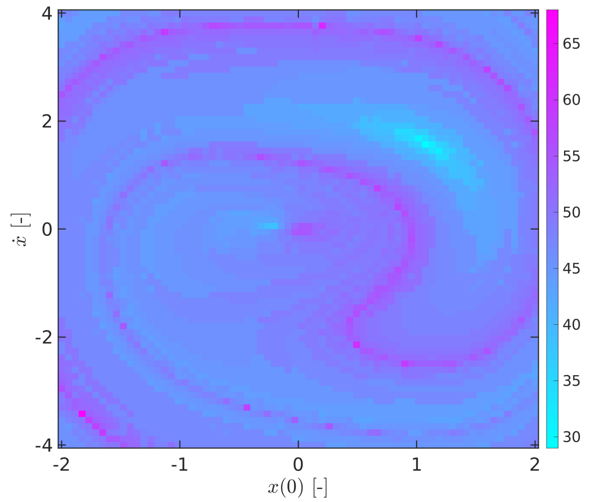

(d) Stage 2 convergence speed in Poincaré periods.

Figure 38: Case 2 results. Despite the larger nonlinearity ratio $\varepsilon = 1$, the elliptic reference reproduces the exact basin throughout the grid.

4.3 Case 3: Unforced Bistable System ($\alpha < 0$)

Case 3 uses $\alpha = -1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0$. The double-well potential has stable equilibria at $x_\pm = \pm 1$ and an unstable equilibrium at the origin. As argued in Section 3.4.3, the sinusoidal reference is structurally inapplicable for in-well cells. The dn elliptic reference is used for cells with $E < E_s$; the cn reference for cross-well cells.

4.3.1 Reference Basin of Attraction

The reference basin for the unforced system is shown in Figure 39. Unlike Cases 1 and 2, where bistability results from frequency-response hysteresis, the basin structure here is governed by the autonomous dynamics. The boundary between the two basins is defined by the stable manifold of the saddle point located at the origin, producing a characteristic interleaved structure that separates the phase plane into two symmetric domains corresponding to the left and right potential wells. The trajectories settle into fixed-point attractors rather than periodic orbits.

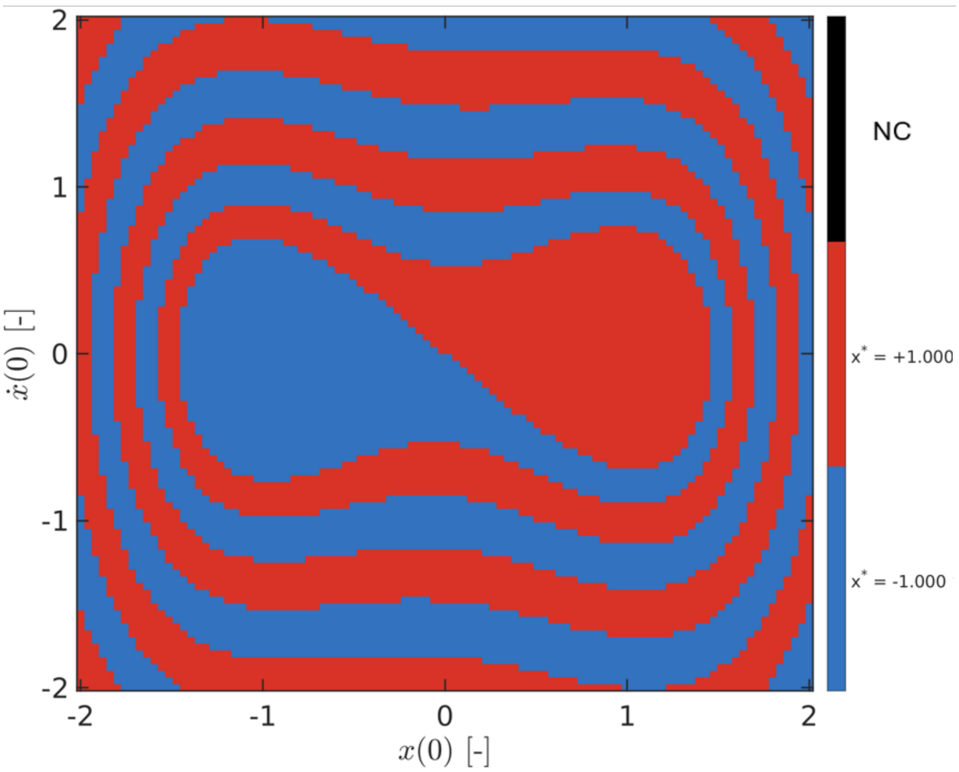

Figure 39: Exact basin for Case 3 (direct integration, no controller).

4.3.2 Controlled Basin of Attraction

The gain maps for Case 3 are computed following Section 3.4.5. Figure 40 shows two clearly distinct regions separated by the separatrix. This is a physically important feature: in-well cells have lower gains because the dn reference oscillates at a lower frequency than the cn reference at the same distance from the origin.

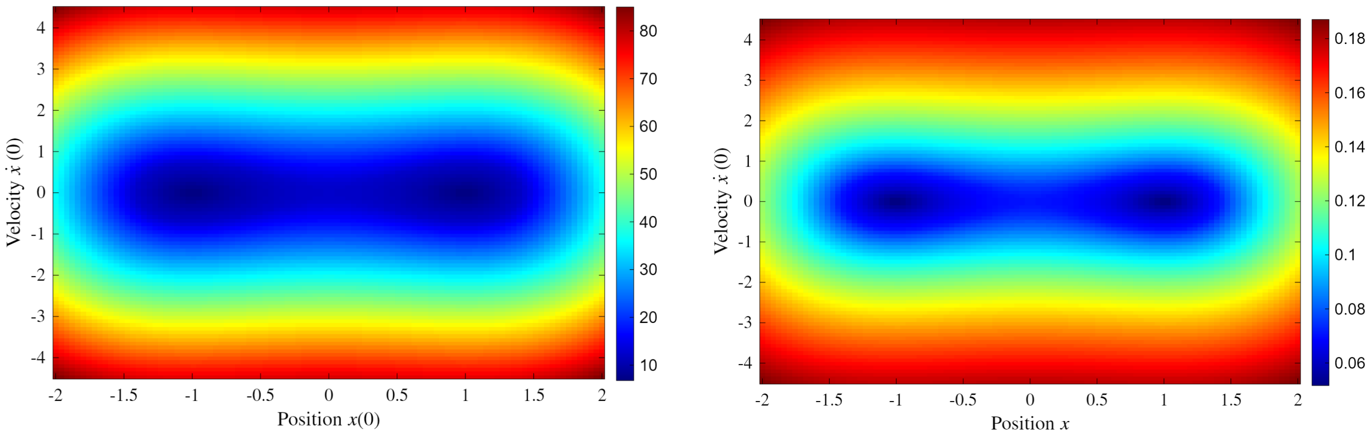

(a) Proportional gain K_p.

(b) Derivative gain K_d.

Figure 40: Scheduled PD gains for Case 3 (elliptic dn/cn reference, $\kappa = 5$, $\zeta = 1/\sqrt{2}$). In-well cells (inside the separatrix) have lower gains than cross-well cells.

Figure 41 gathers the controlled basin, the difference map, and the Stage 2 convergence speed. The two-basin structure is correctly recovered throughout the grid, and all target initial conditions are successfully reached without any Stage 1 timeouts.

Visually, the controlled basin closely matches the reference map. The difference map reveals that a small number of initial conditions converge to the incorrect attractor. These discrepancies are localised near the basin boundary, where the stable manifold of the saddle point creates closely interleaved bands. Minor release errors in this region can easily alter the trajectory and cause a switch in the final attractor.

Regarding the convergence speed, the fastest settling times occur for initial conditions situated very close to the stable equilibrium points. A distinct and separate pattern also emerges at the energy level of the saddle point. Initial conditions initiated near this zero-energy level converge notably faster than states located further down inside the potential wells. This occurs because trajectories starting near the saddle height approach the central region with very low velocity, allowing damping to quickly dissipate their energy so they drop directly into a fixed point. In contrast, initial conditions located within the well interiors but far from the equilibrium undergo a series of underdamped, swinging oscillations, requiring more time to settle. Beyond the saddle energy level, the time required for convergence increases symmetrically with the initial state amplitude, creating a balanced convergence map across the phase space.

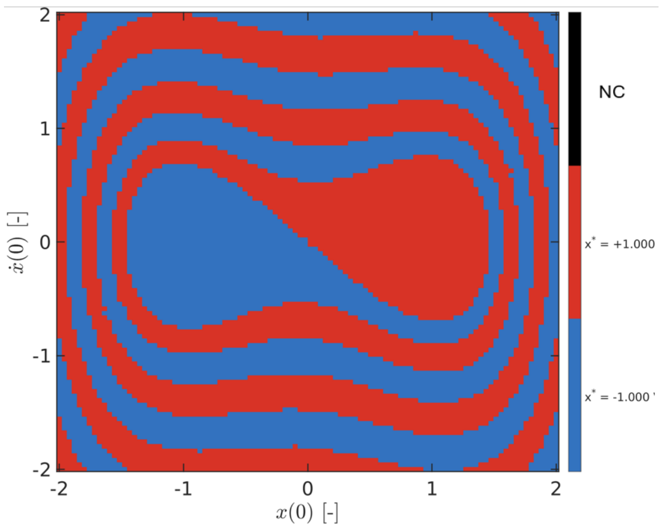

(a) Controlled basin (elliptic dn/cn reference). Same colour convention as Figure 39.

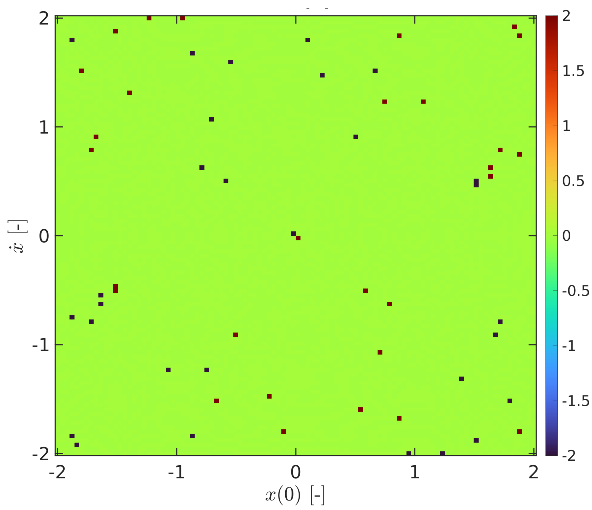

(b) Difference map Δx^*. Discrepancies are localised near the basin boundary.

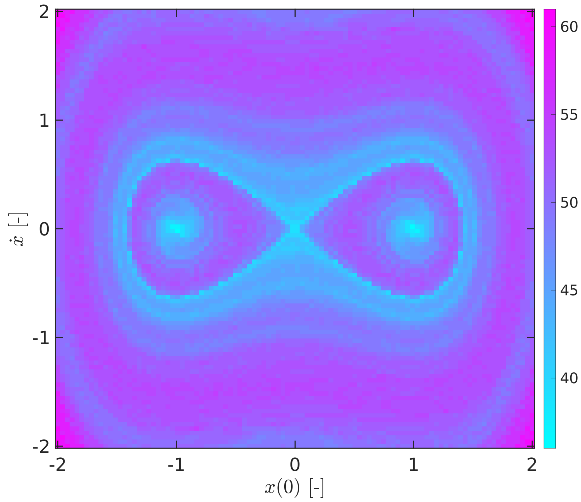

(c) Stage 2 convergence speed in Poincaré periods.

Figure 41: Case 3 results (unforced bistable, elliptic dn/cn reference).

4.4 Basin Evolution Across the Bistable Band

The basin maps of Cases 1 and 2 were each computed at a single operating frequency. The FRC shows that both attractors coexist only over a finite bistable band; as the excitation frequency approaches either fold, one attractor disappears. Because Stage 1 drives the system to a prescribed energy level independently of the forcing frequency, the entire grid sweep can be repeated at a new frequency without modifying the Stage 1 logic: only the Poincaré period $T_f = 2\pi/\omega_f$ changes.

Case 1: basin evolution. Figure 42 shows the controlled basin at two distinct frequencies spanning the bistable band of Case 1.

Close to the lower fold bifurcation, the high-amplitude attractor dominates the phase plane, capturing approximately 80% of the initial conditions. Because critical slowing down is highly prominent near this bifurcation boundary, a significant portion of the trajectories spend extended transients in the phase-space bottleneck, prematurely triggering the convergence criterion before settling onto a primary stable branch. This produces the artificial third-attractor artefact discussed in Section 4.1.3.

Near the upper fold bifurcation, the low-amplitude attractor becomes dominant as the system approaches the transition back to a monostable regime. Manifestations of critical slowing down reappear here as well, producing a similar thin artefact strip along the basin boundary.

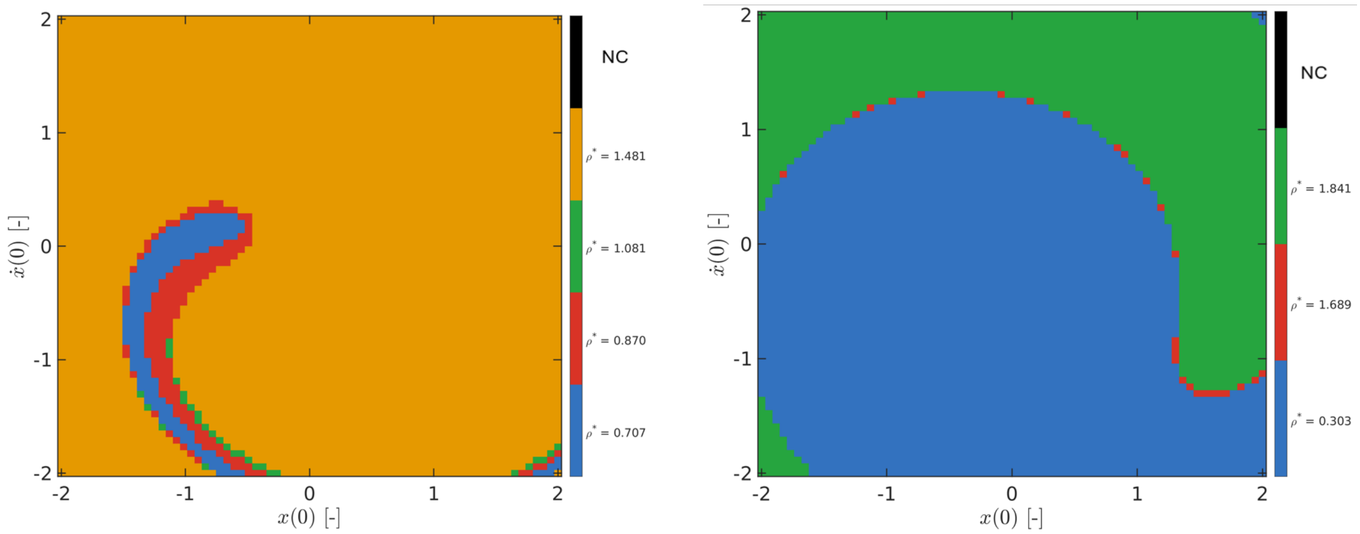

(a) Near the lower fold: the high-amplitude basin dominates.

(b) Near the upper fold: the low-amplitude basin dominates.

Figure 42: Basin evolution across the bistable band for Case 1. The high-amplitude basin erodes continuously as the excitation frequency increases toward the upper fold.

Case 2: basin evolution. Figure 43 illustrates the evolution of the controlled basin across the wider bistable band of Case 2. The shifting dominance follows a clear trend as a function of excitation frequency. Near the lower fold, the low-amplitude attractor dominates. In the middle of the bistable band, the two basins reach an approximate 50/50 proportion. Near the upper fold, the trend reverses and the high-amplitude basin becomes dominant.

Notably, unlike Case 1, the artefacts associated with critical slowing down are completely absent across all three frequency steps. Because the bistable band in Case 2 is significantly wider, none of the chosen operating frequencies sits close enough to the absolute fold bifurcation boundaries to trigger severe critical slowing down during Stage 2.

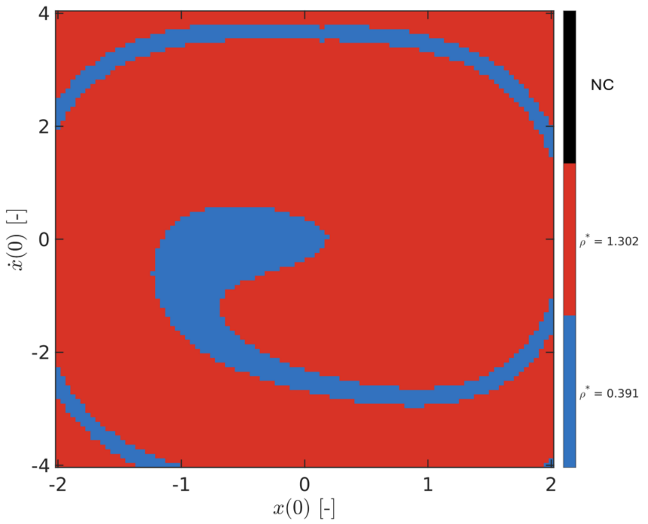

(a) Near the lower fold: low-amplitude basin dominates.

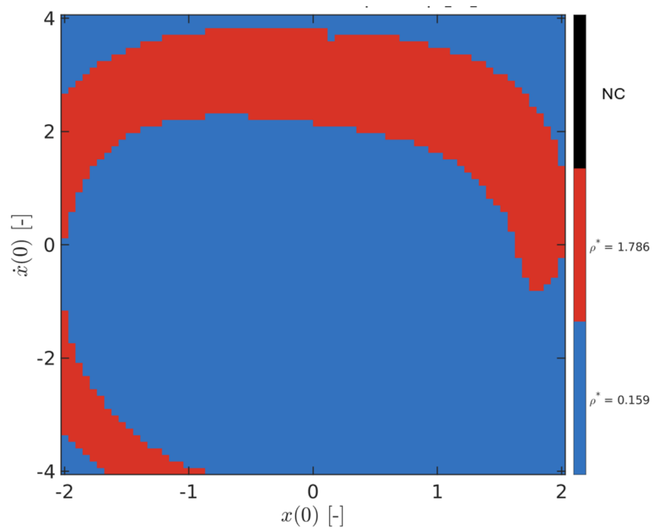

(b) Near the upper fold: high-amplitude basin dominates.

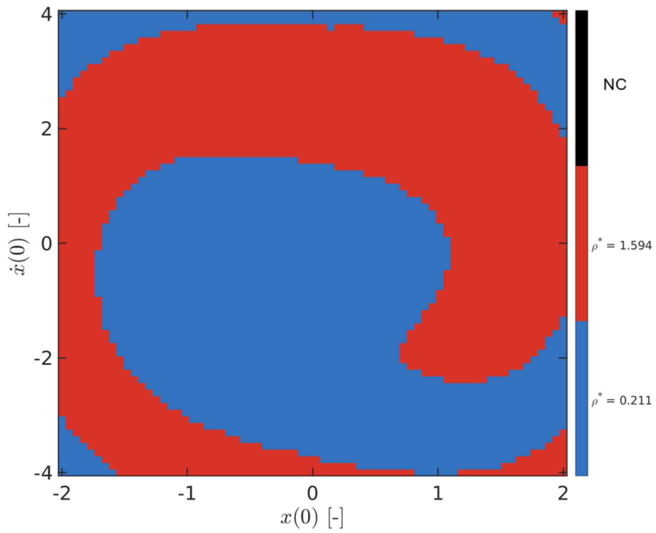

(c) Mid-band: two basins of comparable extent.

Figure 43: Basin evolution across the bistable band for Case 2.

4.5 Parameter Uncertainty Analysis

The preceding cases assumed exact parameter knowledge. In practice, these are uncertain. In the case of the electronic Duffing oscillator considered in Chapter 5, this uncertainty stems from component tolerances and finite potentiometer resolution. This section quantifies the sensitivity of the reconstructed basin using Case 1 as the primary illustration. One physical parameter at a time is perturbed by $\pm 10\%$ (or $\pm 5\%$ for α) while all others are fixed; the reference and gain schedule are recomputed consistently.

4.5.1 Sensitivity to Physical Parameters

Damping coefficient δ. As shown in Figure 44, increasing damping shrinks the high-amplitude basin because the additional dissipation makes it harder to sustain the high-amplitude attractor: cells that would otherwise converge to the high-amplitude orbit are now captured by the low-amplitude one. The effect is moderate and does not alter the overall basin topology.

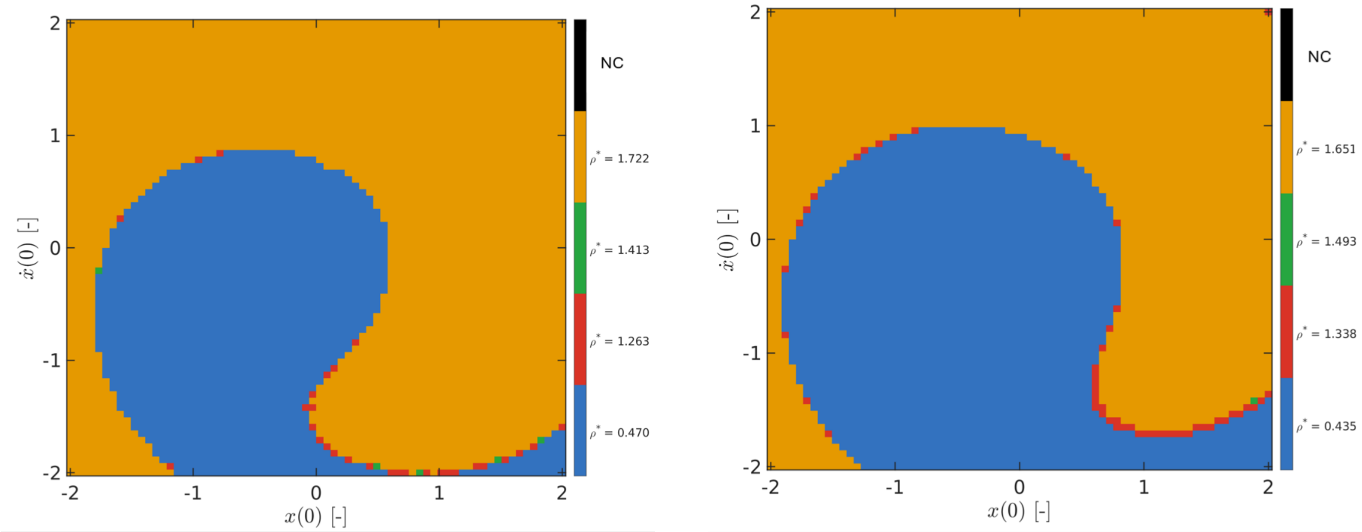

(a) $\delta = 0.9\,\delta_0$ (-10%).

(b) $\delta = 1.1\,\delta_0$ ($+10\%$).

Figure 44: Effect of $\pm 10\%$ perturbation of δ on the controlled basin for Case 1. Nominal map: Figure 29.

Linear stiffness α. The linear stiffness dictates the natural frequency $\omega_0 = \sqrt{\alpha}$ and therefore shifts the FRC along the frequency axis. A $\pm 5\%$ perturbation is applied because this minor variation already moves the fold frequencies noticeably, as shown in Figure 45. Because the system is weakly nonlinear in Case 1, the bistable region is narrow. Consequently, small parameter deviations alter the underlying phase-space topology relative to the nominal operating frequency.

Specifically, a -5% reduction in α shifts the FRC to the left, moving the operating frequency out of the narrow bistable band and into the low-amplitude monostable region. A $+5\%$ increase shifts it to the right, pushing the operating point deeper into the bistability zone and closer to the upper fold. This vulnerability to minor parameter variations makes α the dominant sensitivity parameter for Case 1.

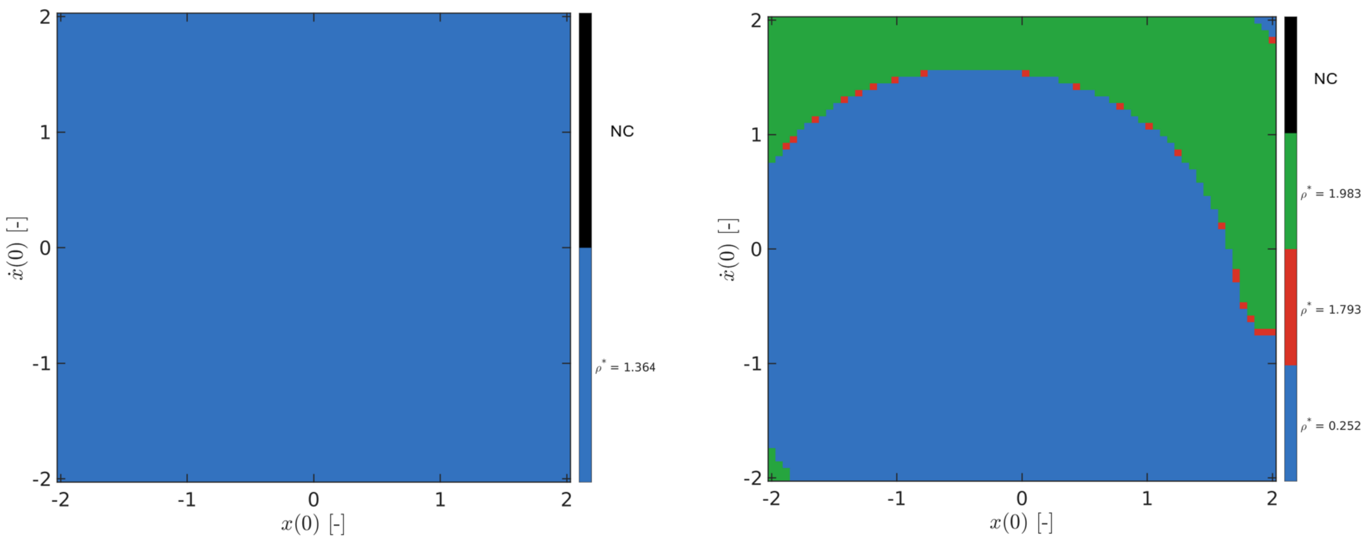

(a) $\alpha = 0.95\,\alpha_0$ (-5%).

(b) $\alpha = 1.05\,\alpha_0$ ($+5\%$).

Figure 45: Effect of $\pm 5\%$ perturbation of α on the controlled basin for Case 1. The basin deformation is already substantial at this perturbation level.

Cubic nonlinearity β. Figure 46 shows the impact of a $\pm 10\%$ perturbation of β. Variations in this parameter modify the bending of the FRC and thus the width of the bistable region. Reducing β by 10% narrows the bistable band, moving the upper fold closer to the nominal operating frequency and allowing the low-amplitude basin to expand. Increasing β by 10% widens the bistable band, pushing the upper fold further away and allowing the high-amplitude basin to occupy a larger portion of the phase plane. The system exhibits lower sensitivity to variations in β than to variations in α.

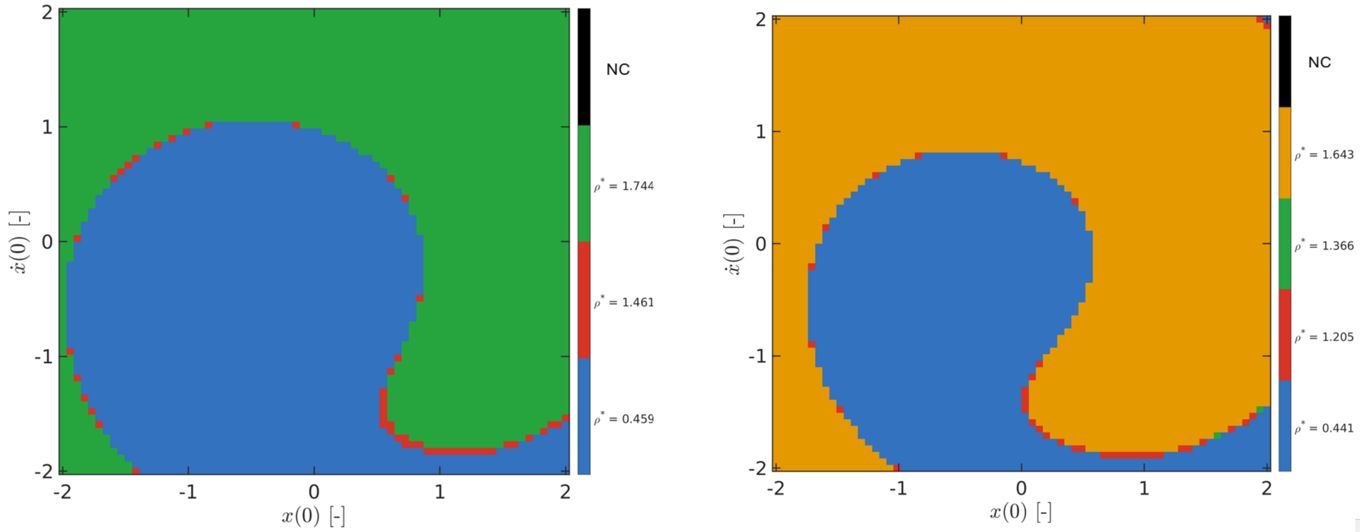

(a) $\beta = 0.9\,\beta_0$ (-10%).

(b) $\beta = 1.1\,\beta_0$ ($+10\%$).

Figure 46: Effect of $\pm10\%$ perturbation of β on the controlled basin for Case 1. The deformation is moderate compared to that induced by α.

Forcing amplitude γ. Figure 47 illustrates the effects of perturbing γ. A reduction in γ shrinks the bistable band and shifts the balance between the two basins, allowing the low-amplitude attractor to occupy a larger portion of the phase plane. An increase widens the bistable region and expands the high-amplitude basin. Variations in the forcing amplitude also directly alter the steady-state response amplitudes of the stable branches, shifting the basin centroids within the phase space.

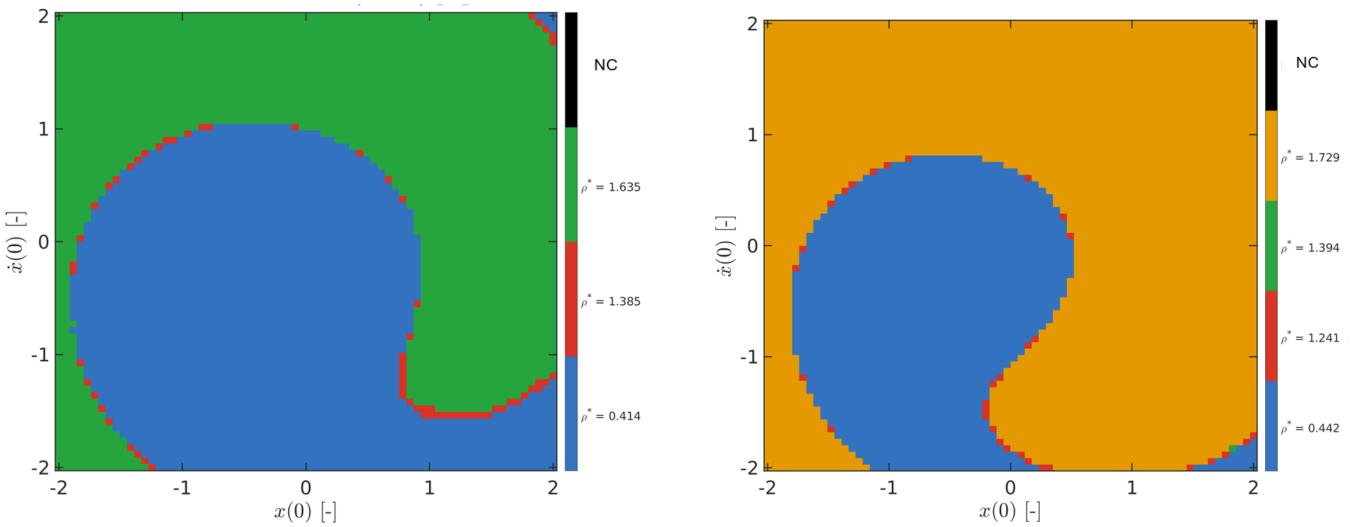

(a) $\gamma = 0.9\,\gamma_0$ (-10%).

(b) $\gamma = 1.1\,\gamma_0$ ($+10\%$).

Figure 47: Effect of $\pm10\%$ perturbation of γ on the controlled basin for Case 1. The basin boundary migrates consistently with the shift of the attractors on the FRC.

4.5.2 Sensitivity to Algorithm Parameters

Poincaré convergence tolerance τ_{pc}. Figure 48 shows the influence of the convergence tolerance on the reconstructed basin. A loose tolerance ($\tau_{\mathrm{pc}} = 0.1$) causes premature steady-state declarations, producing salt-and-pepper noise across the basin interiors. Tightening the threshold to $\tau_{\mathrm{pc}} = 0.01$ eliminates the majority of these artefacts without altering the primary basin boundaries. Further tightening to $\tau_{\text{pc}} = 0.001$ yields marginal differences at the current grid resolution; its primary effect is the resolution of a few remaining initial conditions along the boundary that were previously captured by the critical slowing down bottleneck. Because these localised corrections do not change the macro-topology of the basins, and because this tolerance level is also feasible in a noisy experiment, $\tau_{\text{pc}} = 0.01$ is adopted as the standard value.

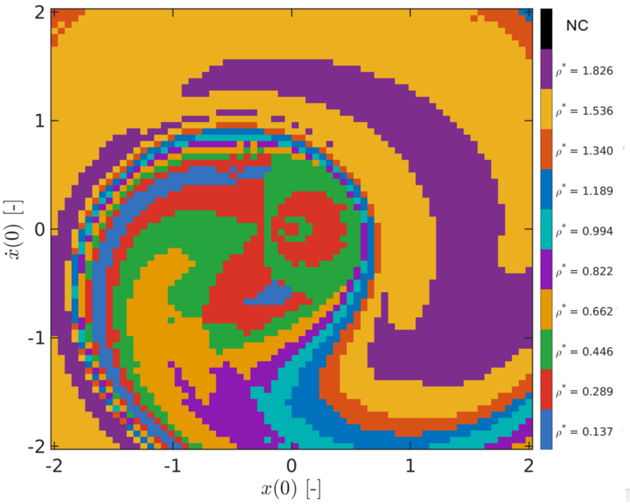

(a) $\tau_{\text{pc}} = 0.1$: premature convergence produces salt-and-pepper noise.

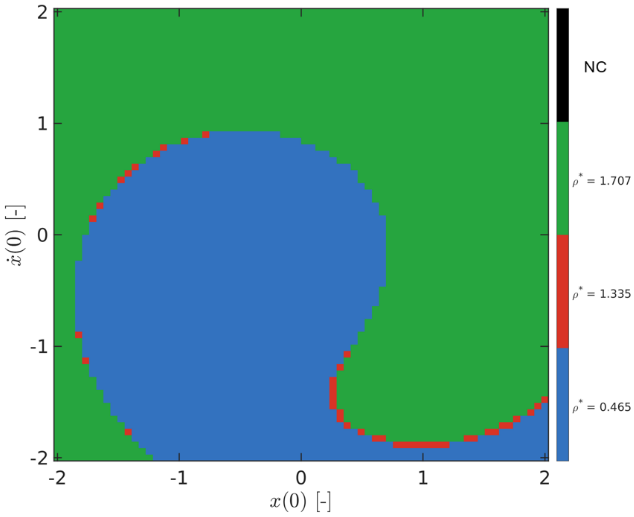

(b) $\tau_{\text{pc}} = 0.01$: noise eliminated; basin well resolved.

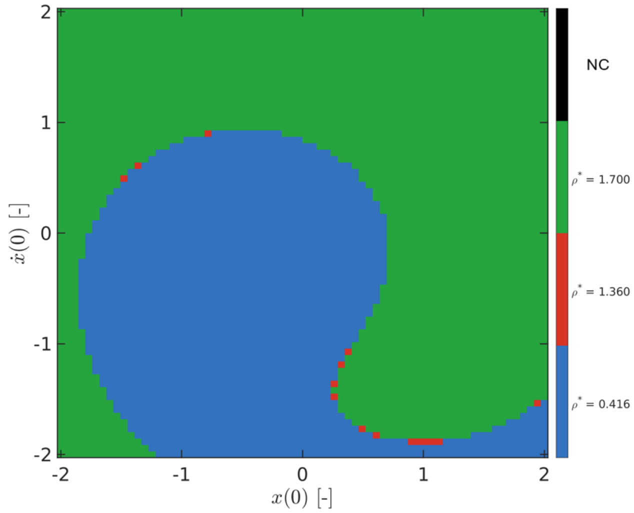

(c) $\tau_{\text{pc}} = 0.001$: minor differences compared to the 0.01 case.

Figure 48: Sensitivity to τ_{pc} (Case 1, $n_{\text{stable}} = 20$). The value $\tau_{\text{pc}} = 0.01$ is adopted.

Stable-count threshold n_{stable}. Figure 49 illustrates the sensitivity to the stable-count threshold n_{stable} under the nominal tolerance ($\tau_{\text{pc}} = 0.01$). This parameter acts as the primary safeguard against premature convergence declarations caused by critical slowing down. Near the lower fold bifurcation, trajectories entering the phase-space bottleneck exhibit extended transients that mimic steady state over multiple periods.

When a low threshold is used ($n_{\text{stable}} = 2$), the algorithm interprets this temporary stabilisation as final convergence, producing a large artificial region attributed to the unstable branch. Increasing the threshold to $n_{\text{stable}} = 4$ filters out some artefacts, though a distinct band persists along the basin boundary. Raising the requirement to $n_{\text{stable}} = 20$ monitors the system long enough to outlast the majority of these extended transients, resolving the primary structure of the true boundary. Further increasing to $n_{\text{stable}} = 40$ leads to near-total disappearance of the remaining artefacts, confirming theoretical convergence behaviour. Because the macro-topology is already sufficiently captured at $n_{\text{stable}} = 20$ and a higher threshold increases the overall computational cost of the sweep, $n_{\text{stable}} = 20$ is selected as a practical compromise.

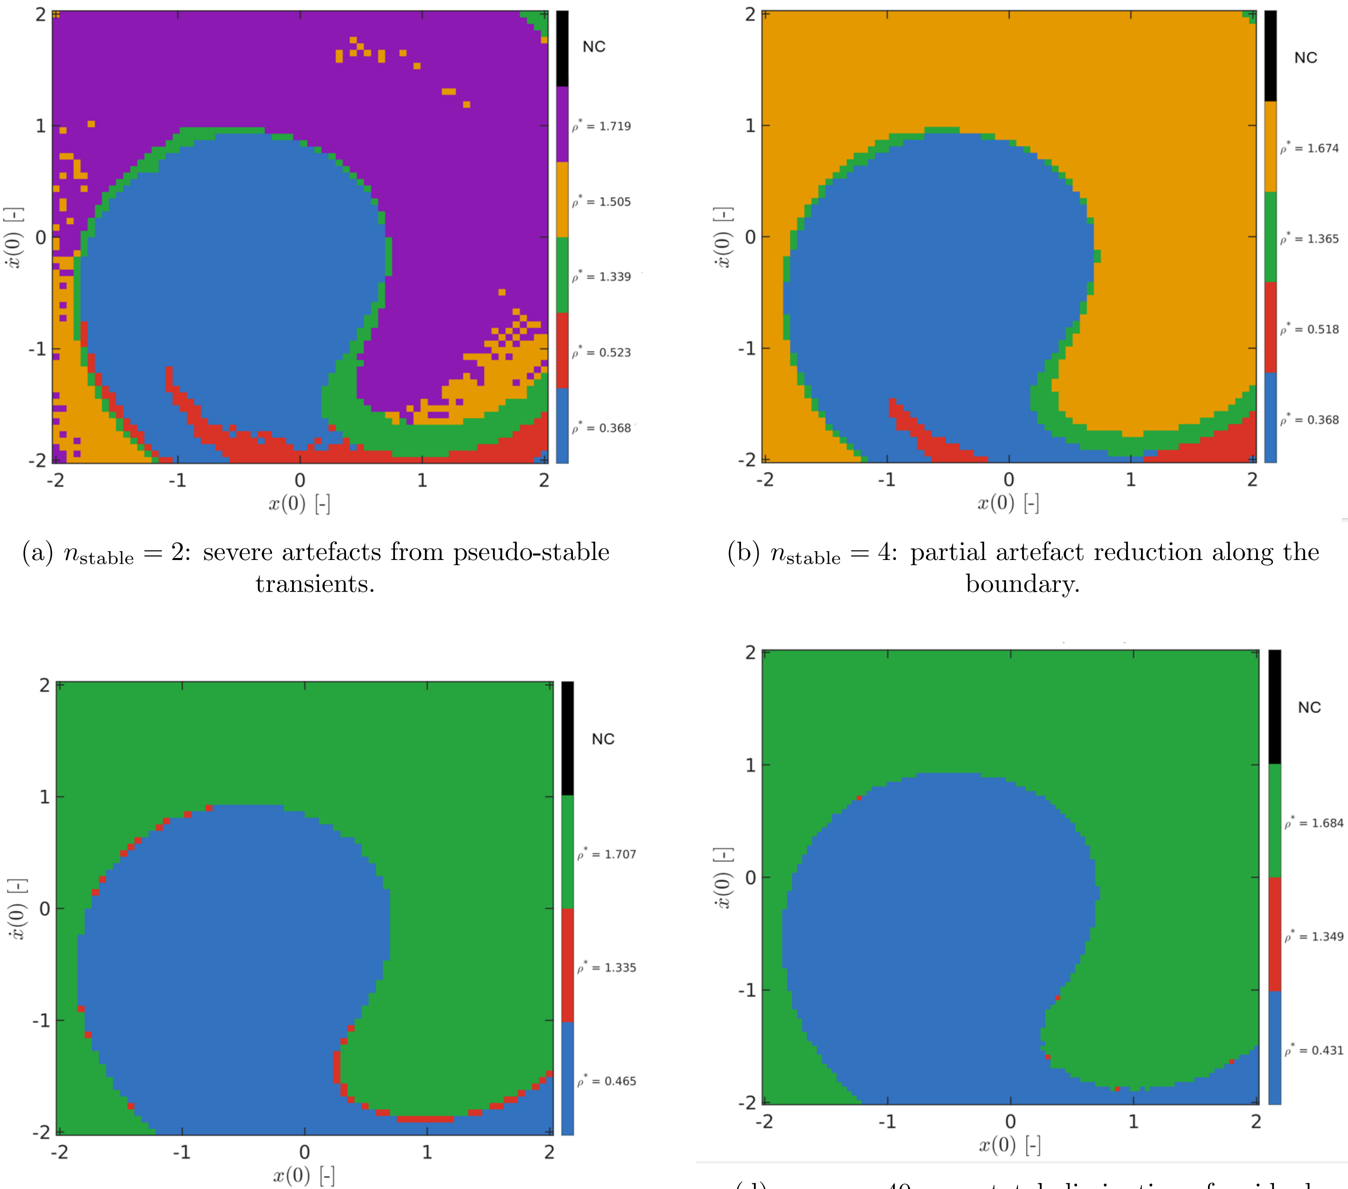

(a) $n_{\text{stable}} = 2$: severe artefacts from pseudo-stable transients.

(b) $n_{\text{stable}} = 4$: partial artefact reduction along the boundary.

(c) $n_{\text{stable}} = 20$: core boundary resolved; adopted value.

(d) $n_{\text{stable}} = 40$: near-total elimination of residual artefacts.

Figure 49: Sensitivity to n_{stable} (Case 1, $\tau_{\text{pc}} = 0.01$).

4.6 Case 4: Highly Forced System: Subharmonics and Isolated Response Branches

4.6.1 Motivation and System Parameters

Case 4 uses $\alpha = 1$, $\beta = 1$, $\delta = 0.05$, $\gamma = 3$, inspired by [16]. The very low damping and large forcing amplitude place the system well outside the weakly nonlinear regime, giving rise simultaneously to fundamental bistability and 1:2 and 1:3 subharmonic resonances.

4.6.2 Frequency-Response Curve

Figure 50 shows the FRC for Case 4, which exhibits a broad multistability region extending from $\omega_f \approx 2$ to 7.5. Within this band the system undergoes several topological transitions. Near the lower fold bifurcation, a standard bistable regime is initially observed. Between $\omega_f = 2.5$ and 3, a tristable region emerges due to the 1:2 subharmonic isola. Another tristable regime occurs from $\omega_f = 4$ up to the upper fold, driven by the 1:3 subharmonic branch. A narrow window also exists where the high-frequency termination of the 1:2 isola overlaps with the initiation of the 1:3 isola.

To systematically analyse these coexisting states, the controlled basin investigations are carried out at four distinct operating frequencies: the initial bistable zone, the tristable region containing the 1:2 isola, the narrow overlap window where both subharmonic structures coexist, and the final tristable region governed by the 1:3 subharmonic branch.

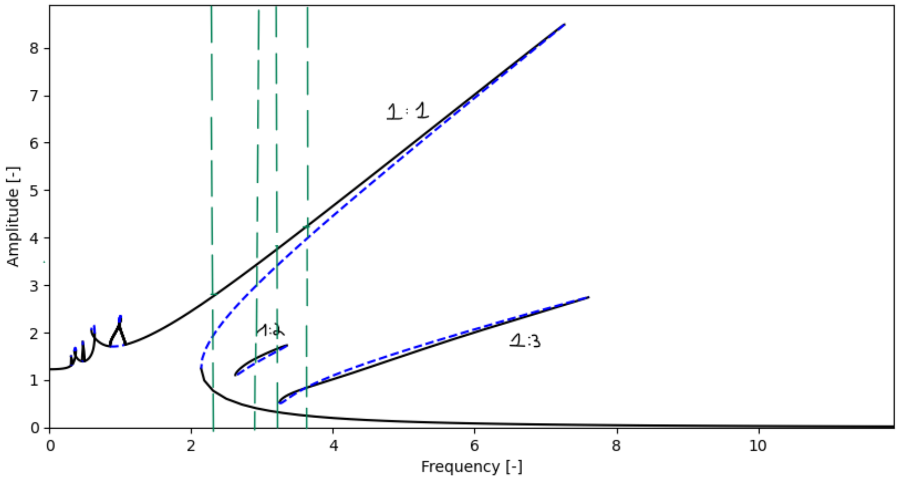

Figure 50: FRC for Case 4 ($\alpha = 1$, $\beta = 1$, $\delta = 0.05$, $\gamma = 3$). The disconnected branches correspond to the 1:2 and 1:3 subharmonic resonances.

4.6.3 Basin Maps at Four Representative Frequencies

Figure 51 shows the controlled basin at the four forcing frequencies described above.

In Figure 51(a), chosen within the initial bistable zone, only the low- and high-amplitude attractors coexist. The high-amplitude basin dominates the majority of the grid.

In Figure 51(b), operating within the 1:2 subharmonic regime, the basin of the isola is visible in green. The initial conditions converging to this isola do not form a single contiguous domain; instead, they appear as distinct localised regions embedded within the low-amplitude basin. Topologically, these manifest as four symmetrical elongated lobes circulating around the origin. A fourth artificial basin also appears with an amplitude matching that of the isola, caused by extended transients near the true isola basin boundaries.

In Figure 51(c), the operating frequency is inside the narrow overlap window. The primary low- and high-amplitude basins are visible alongside the 1:3 subharmonic branch, whose three symmetrical tendril structures are captured entirely within the low-amplitude basin. The residual green points near the decaying 1:2 isola boundary are transient artefacts analogous to those discussed above.

In Figure 51(d), the phase space resolves into exactly three primary domains with no residual artefacts: the low-amplitude basin, the high-amplitude basin, and the fully developed 1:3 subharmonic region.

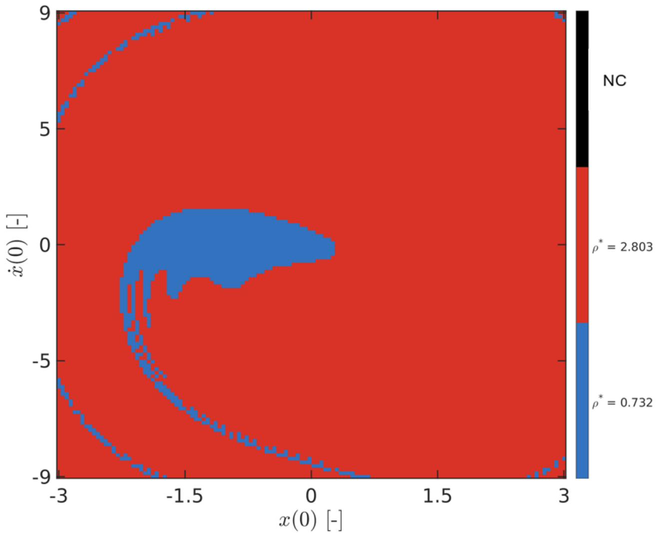

(a) Low frequency: fundamental bistability only.

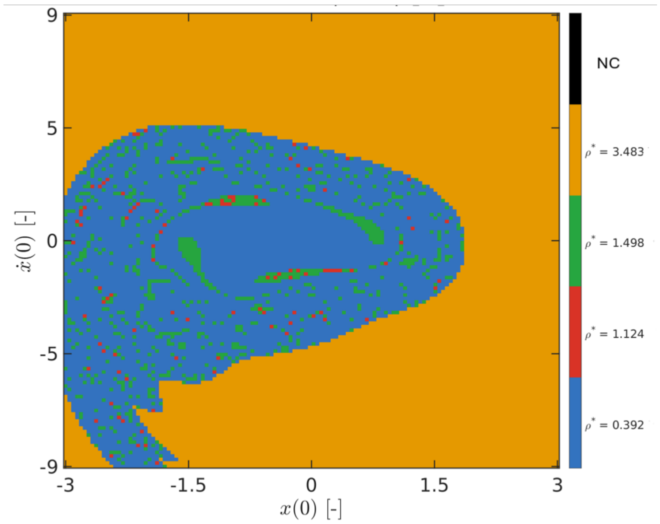

(b) Lower mid-band: 1:2 subharmonic basin (green) embedded within the low-amplitude basin.

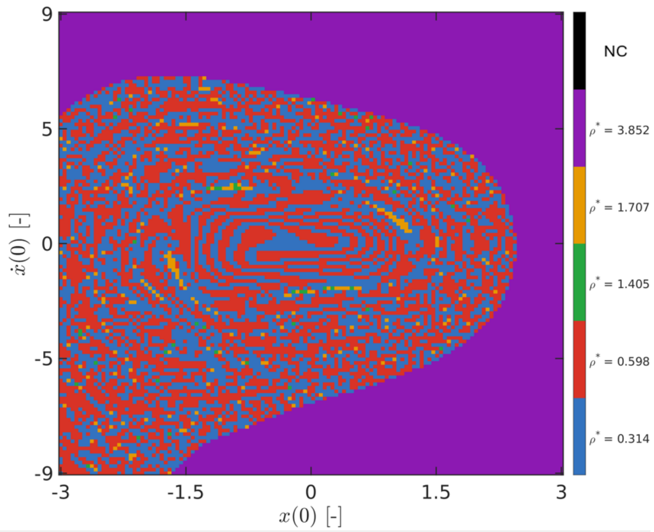

(c) Overlap window: 1:3 subharmonic basin (red) alongside residual 1:2 isola boundary points(yellow).

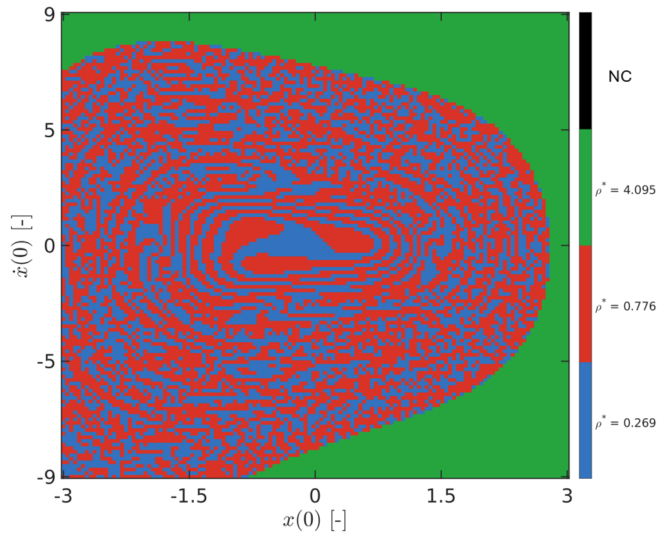

(d) High frequency: pure tristable regime with fully developed 1:3 subharmonic region (red).

Figure 51: Controlled basin for Case 4 at four forcing frequencies.

4.7 Backfilling Basin Maps

The backfilling procedure described in Section 3.6.2 is applied to all four configurations. The motivation is to reduce the total number of active two-stage integrations required. Rather than running Stage 1 and Stage 2 for every cell independently, only a subset of cells are actively integrated; the remaining cells are assigned an attractor label by post-processing the stroboscopic states visited during the Stage 2 trajectories of the active cells.

Concretely, active initial conditions were distributed in a frame-like region surrounding the rectangular grid considered in the preceding sections. This frame comprised a total number of initial conditions equal to 30% of the full grid. Only these frame cells were actively integrated. The stroboscopic Poincaré states recorded during their Stage 2 evolutions were then used to retroactively assign labels to cells inside the rectangular region, without running additional Stage 1 or Stage 2 procedures for those cells. The remaining cells inside the rectangular region that were not reached by any stroboscopic trajectory are left unclassified.

For each configuration, two maps are presented: the reconstructed basin (adopting the same colour convention as the preceding sections) and the classification status map. In the status maps, turquoise cells were classified by direct integration (status 1), yellow cells were labelled by the backfill procedure (status 2), and dark blue cells remain unclassified. Figures 52–55 present the results for all four cases.

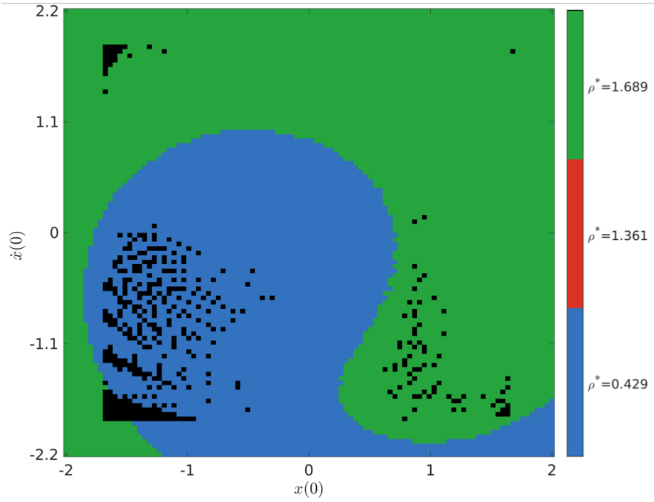

(a) Reconstructed basin from 30% active integration.

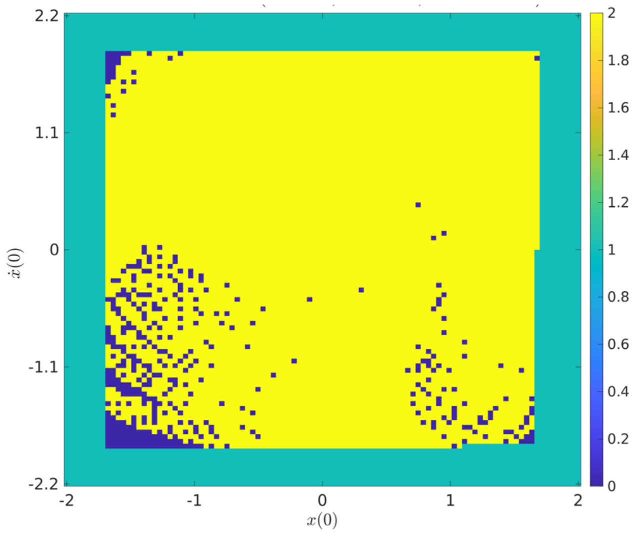

(b) Classification status map (turquoise: direct, status 1; yellow: backfilled, status 2; dark blue: unclassified).

Figure 52: Case 1 backfill result ($\alpha = 1$, $\beta = 0.052$, $\delta = 0.016$, $\gamma = 0.039$, $\nu = 1.046$).

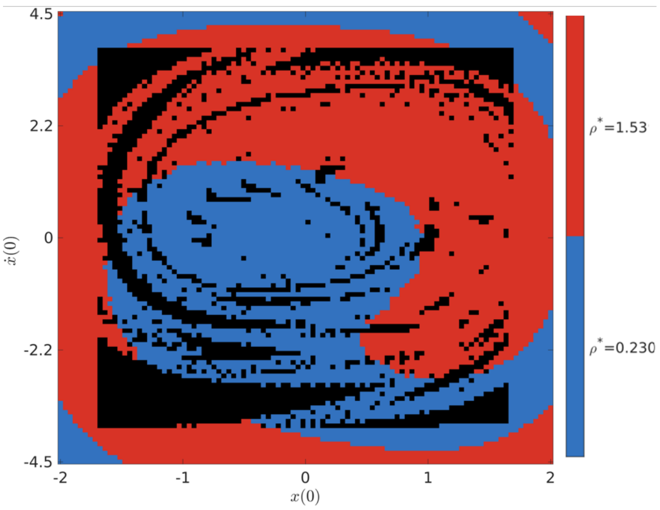

(a) Reconstructed basin from 30% active integration.

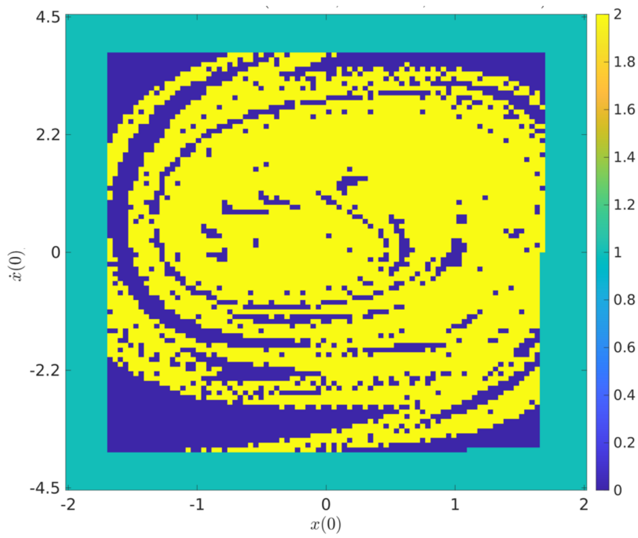

(b) Classification status map.

Figure 53: Case 2 backfill result ($\alpha = 1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0.35$, $\nu = 1.6$).

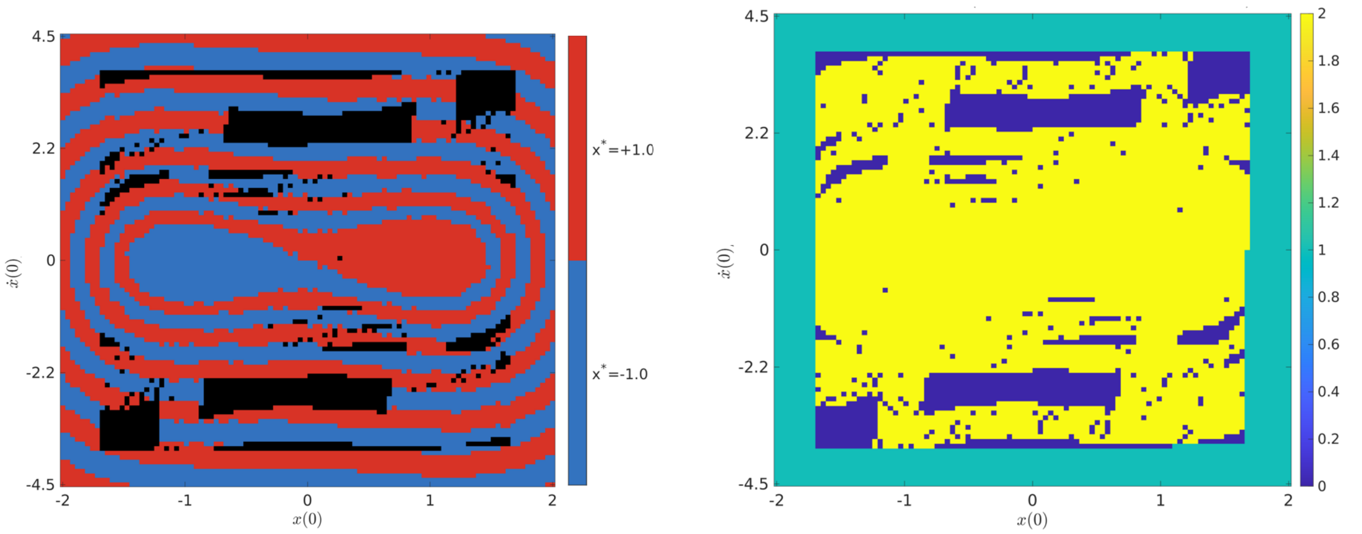

(a) Reconstructed basin from 30% active integration.

(b) Classification status map.

Figure 54: Case 3 backfill result ($\alpha = -1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0$).

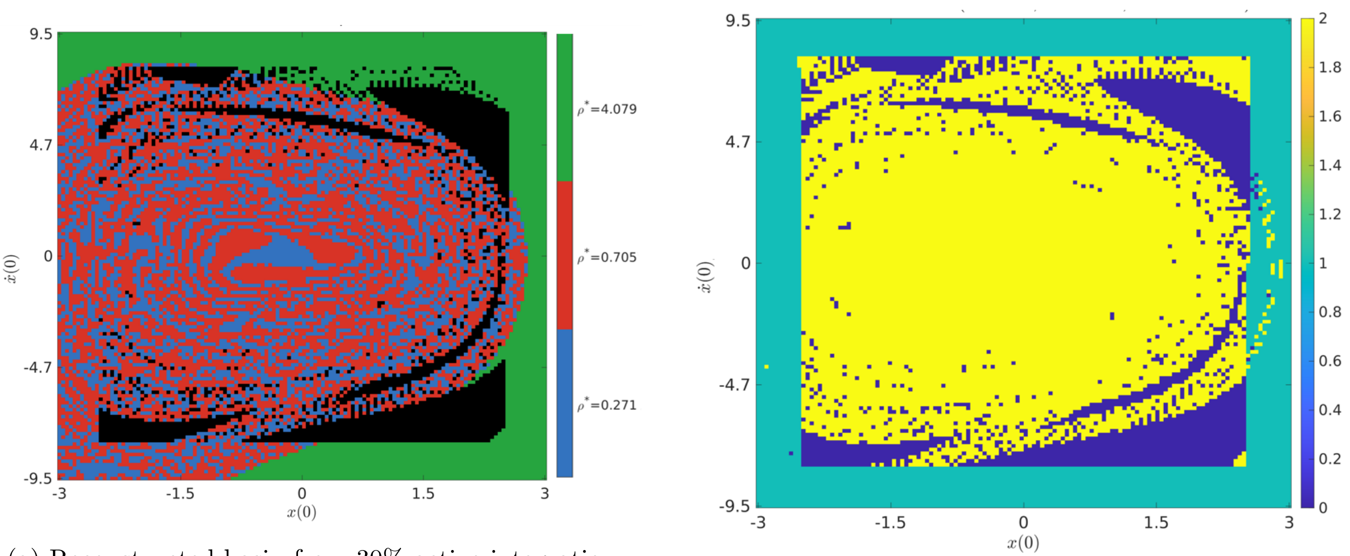

(a) Reconstructed basin from 30% active integration. The subharmonic basins are recovered.

(b) Classification status map.

Figure 55: Case 4 backfill result ($\alpha = 1$, $\beta = 1$, $\delta = 0.05$, $\gamma = 3$). The backfill procedure correctly propagates labels through the subharmonic basin regions without misclassification.

The backfilling algorithm achieves a faithful replication of the true basin topologies across all four cases. Even when limited to an active integration budget of 30% of the grid, the procedure successfully resolves macro-structures, fine boundaries, and localised subharmonic features without introducing misclassification artefacts. However, in Case 4, a small proportion of cells within the 30% budget are backfill-labelled rather than directly integrated, producing no visible degradation in accuracy. The efficiency of the approach depends on the system's damping level: a lightly damped system has long transients, so each Stage 2 trajectory visits many stroboscopic cells before declaring convergence, and the backfill propagates labels across a wide area of the grid. A heavily damped system converges in fewer periods, the visited list is shorter, and the backfill yield is correspondingly lower.

5 Experimental Validation

The numerical validation established that the algorithm is theoretically sound across four qualitatively distinct configurations. This chapter addresses the more demanding question of whether the same strategy can be executed on a physical dynamical system, subject to measurement noise, analogue component tolerances, DAC(Digital-to-Analog Converter) output saturation, and fixed-step real-time computational constraints.

The chapter follows the same case structure as Section 4, so that every experimental result has an unambiguous numerical counterpart. For each case the experimental basin is presented alongside the results based on a numerical model of the oscillator, which serves as the reference for comparison throughout this chapter, enabling a cell-by-cell assessment.

5.1 Experimental Platform

5.1.1 The Electronic Duffing Oscillator

The oscillator is built around three operational amplifiers, two analogue multipliers, and a network of resistors and capacitors. Three potentiometers with division ratios $p_1, p_2, p_3 \in [0, 1]$ allow independent tuning of damping, linear stiffness, and cubic nonlinearity. The output voltage V_{out} represents displacement; circuit analysis [14] yields

$$C_1C_2R_4R_5\,\ddot{V}_{\text{out}} + \frac{p_1C_2R_4R_5}{R_1}\,\dot{V}_{\text{out}} + \frac{p_2R_4R_7}{R_2R_6}\,V_{\text{out}} + \frac{p_3g_m^2R_4R_7^3}{R_3R_6^3}\,V_{\text{out}}^3 = V_{\text{in}}, \tag{36}$$

which is structurally identical to $m\ddot{x} + c\dot{x} + kx + k_3x^3 = f(t)$ with $m = C_1C_2R_4R_5$, $c = p_1C_2R_4R_5/R_1$, $k = p_2R_4R_7/(R_2R_6)$, and $k_3 = p_3g_m^2R_4R_7^3/(R_3R_6^3)$. The fixed component values give $m = 10^{-4}$ s^2.

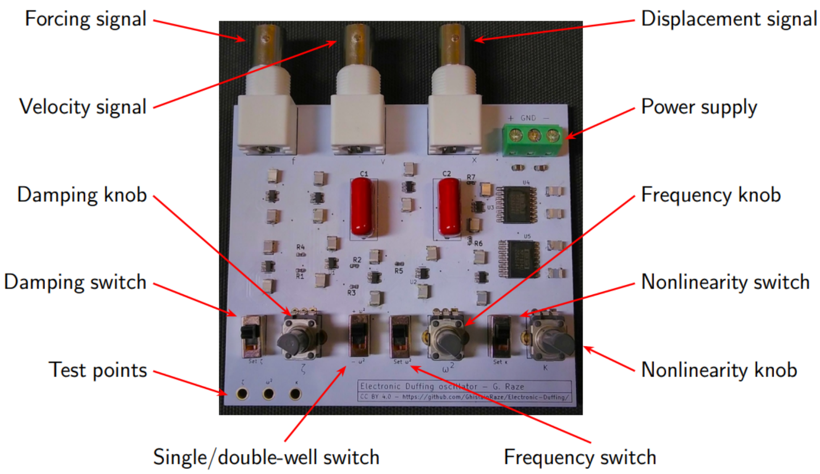

Figure 56: The electronic Duffing oscillator [14].

The three BNC connectors provide access to V_{out} (displacement x, sampled at 10 kHz), the velocity signal $V_1 \approx -C_2R_5\dot{V}_{\text{out}}$ (sign inverted and scaled by 1/1000 to recover $\dot{x}$), and V_{in} (control input from the DAC).

To map a target dimensionless system defined by the parameter set $(\delta, \beta, \gamma, \nu)$ to the physical analogue hardware, the underlying differential equations must be scaled. Defining the natural frequency of the linear system as $\omega_0 = \sqrt{k/m}$, the standard dimensionless variables are introduced through $\tau = \omega_0 t$, with $\dot{x} = \omega_0 x'$ and $\ddot{x} = \omega_0^2 x''$, where primes denote differentiation with respect to τ. Substituting into the physical governing equation and dividing through by the linear stiffness k gives

$$x'' + \frac{c}{m\omega_0}x' + x + \frac{k_3}{k}x^3 = \frac{F}{k}\sin\left(\frac{\Omega}{\omega_0}\tau\right), \tag{37}$$

which, upon identifying the standard dimensionless groups

$$\delta = \frac{c}{m\omega_0}, \quad \beta = \frac{k_3}{k}, \quad \gamma = \frac{F}{k}, \quad \nu = \frac{\Omega}{\omega_0}, \tag{38}$$

simplifies directly to the canonical dimensionless Duffing equation

$$x'' + \delta x' + x + \beta x^3 = \gamma \sin(\nu\tau). \tag{39}$$

Given this equivalence, the physical system parameters required for experimental calibration are uniquely determined from the chosen dimensionless parameters and the natural frequency ω_0 via the inverse mapping

$$k = m\omega_0^2, \quad c = \delta m\omega_0, \quad k_3 = \beta m\omega_0^2, \quad F = \gamma m\omega_0^2, \quad \Omega = \nu\omega_0. \tag{40}$$

5.1.2 Real-Time Implementation: dSPACE MicroLabBox

The tracking controller, convergence detection algorithm, and automated basin stepping are executed in real time on a dSPACE MicroLabBox at a fixed sampling interval of $\Delta t = 10^{-4}$ s. The underlying control model contains three main functional blocks. The first is the reference generator, which evaluates the precomputed state references using dynamic lookup tables and applies the raised-cosine transition ramp defined in equation (26) to smoothly bridge consecutive initial conditions. The second is the tracking controller, which computes the real-time control effort via the PD law using discrete, cell-specific gains tuned to local phase-space dynamics. The third is the convergence detector, which maintains a persistent Poincaré sample history within the real-time processor memory to evaluate steady-state synchronisation.

To enable offline reconstruction of the experimental basins of attraction, four essential signals are logged during the grid stepping. The system records the index of the active grid cell under test to map spatial coordinates, a binary state flag that indicates when the tracking phase terminates and the system is released into free or forced integration, a real-time convergence indicator that flags when a trajectory has successfully matched a steady-state attractor, and a stroboscopic timing signal that triggers at the predefined Poincaré phases.

5.1.3 Physical Non-Idealities and Their Mitigation

A DC offset arising from analogue multiplier bias produces a non-zero output when $V_{\text{in}} = 0$. To mitigate this effect, a constant corrective voltage is injected into the external force input to cancel the electrical offset before each experimental run. After this calibration step, the residual offset is verified to remain below 5 mV. The DAC output is hard-limited to ± 10 V.

Besides these mitigated non-idealities, the electronic oscillator closely behaves like an ideal one, allowing direct comparison with its numerical counterpart.

5.2 Case 1: Weakly Nonlinear Forced Bistable System

Three experimental stepping sequences are performed to evaluate the framework across different operating regimes: a sinusoidal reference at the nominal frequency, an elliptic reference at the same frequency, and an elliptic reference near the upper fold bifurcation.

5.2.1 Experimental Validation with Sinusoidal Reference

The experimental characterisation is conducted using a spiral stepping sequence of initial conditions distributed across the phase-plane grid. The normalised release error, defined as the value of $d(t)$ in equation (30) at the moment the controller is switched off, is used as the primary metric for Stage 1 performance. The results indicate that the controlled tracking phase successfully guides every initial condition into its targeted phase-space cell prior to release: the release error consistently remains below the threshold of 0.5, satisfying the cell-entry criterion of Section 3.4.6. The steady-state response amplitudes recorded after release show excellent agreement with the theoretical values derived from the frequency-response curve.

An overlay of the experimental data points with the numerical basin reference reveals a high degree of topological correspondence, as shown in Figure 57. Minor discrepancies are observed primarily along the basin boundary, where small release errors can push a trajectory toward the wrong attractor. As demonstrated in Section 4.5, even small deviations between the physical analogue parameters and the nominal values can visibly shift the basin boundary, which explains these localised variations.

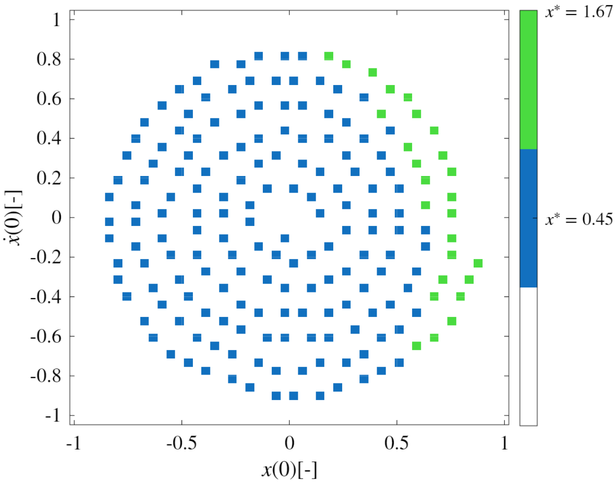

(a) Experimental spiral BoA (sinusoidal reference).

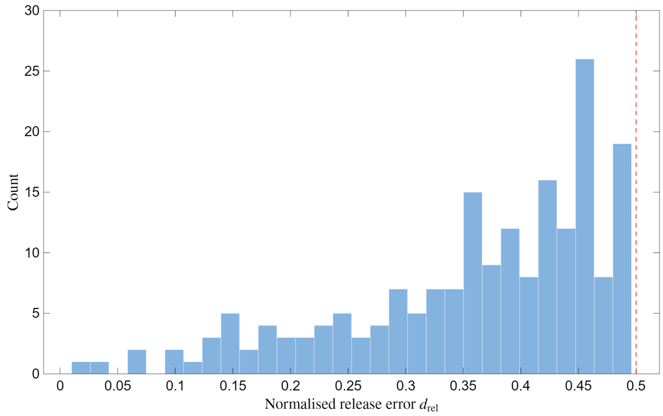

(b) Normalised release error distribution. All completed cells satisfy the threshold $\eta_{\mathrm{rel}} < 0.5$.

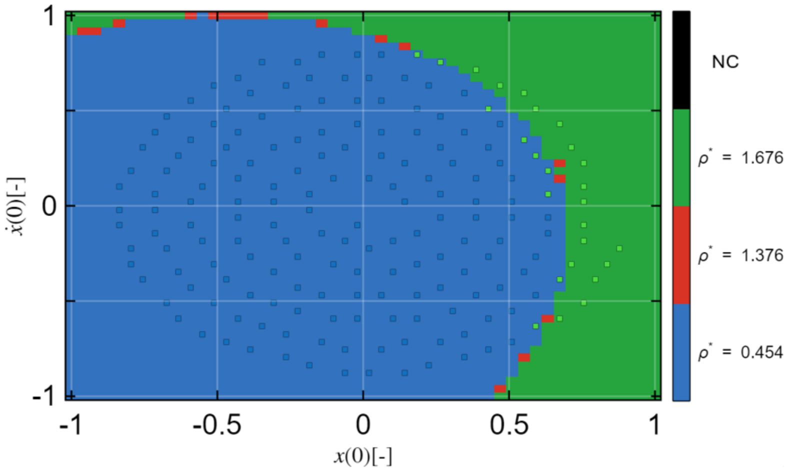

(c) Experimental BoA overlaid on the numerical reference of Figure 29.

Figure 57: Case 1, sinusoidal reference, nominal frequency $\nu_{\mathrm{op}} = 1.046$.

5.2.2 Experimental Validation with Elliptic Reference

When the phase-space grid is extended to larger amplitudes, the sinusoidal reference encounters tracking performance limitations and fails to reliably guide trajectories to the targeted initial conditions, consistently with the numerical observations of Section 4.1.5. The elliptic reference is therefore used instead. An experimental stepping sequence of 200 initial conditions is performed. The steady-state amplitudes identified after release for both branches show excellent quantitative agreement with the frequency-response curve at the operating frequency. The normalised release error remains below the threshold for the vast majority of cells. The overlay shown in Figure 58 demonstrates that the experimental classifications are virtually indistinguishable from the numerical reference, confirming that the elliptic reference corrects the tracking failures observed with the sinusoidal one.

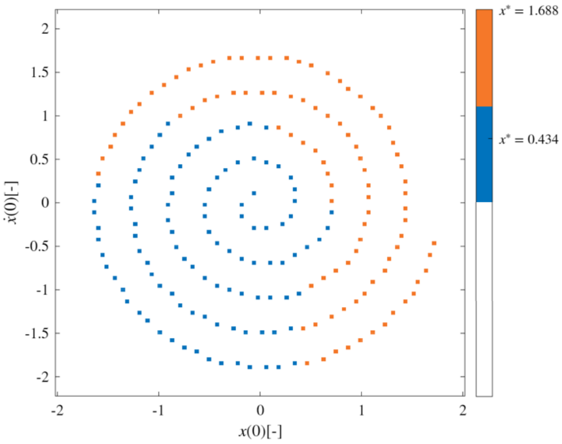

(a) Experimental spiral BoA (elliptic reference).

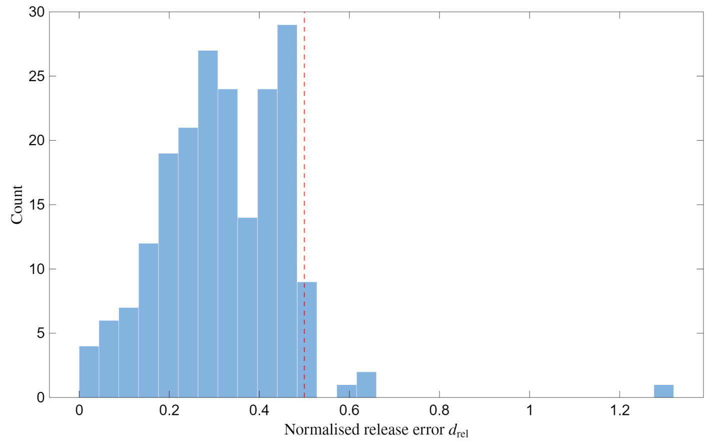

(b) Normalised release error distribution.

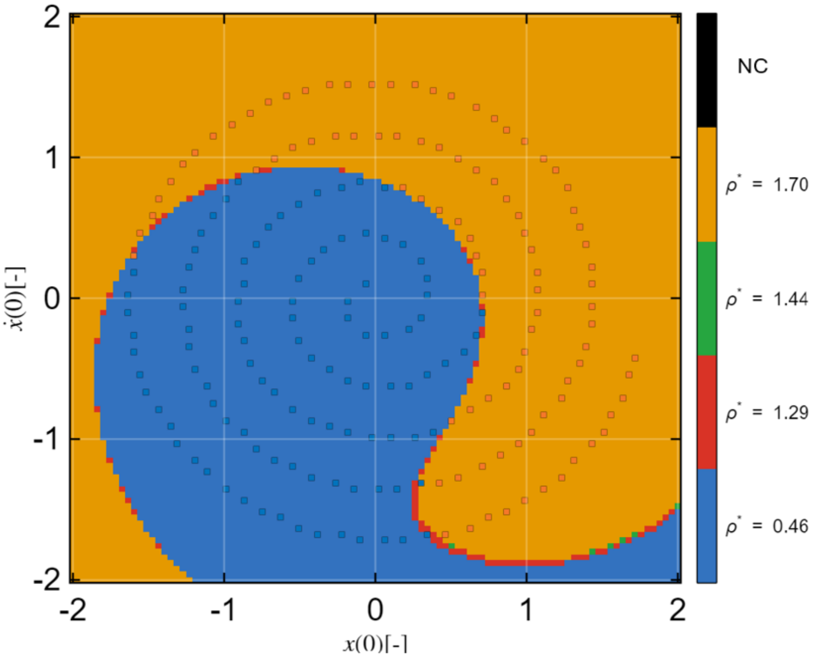

(c) Overlay on the numerical reference.

Figure 58: Case 1, elliptic reference, nominal frequency $\nu_{\mathrm{op}} = 1.046$.

5.2.3 Experimental Validation at a Higher Frequency

Operating closer to the upper fold bifurcation tests the framework under increased phase-space sensitivity and narrowing basin margins. Because the tracking controller is active only prior to release, its operation is unchanged at this frequency; this stepping sequence primarily validates the steady-state detection and basin mapping routines in a more demanding geometric configuration. The normalised release error remains below the threshold throughout, and the overlay with the corresponding numerical frequency-sweep frame of Figure 42 demonstrates excellent agreement, as shown in Figure 59. The erosion of the high-amplitude basin as the operating frequency approaches the upper fold is correctly reproduced experimentally.

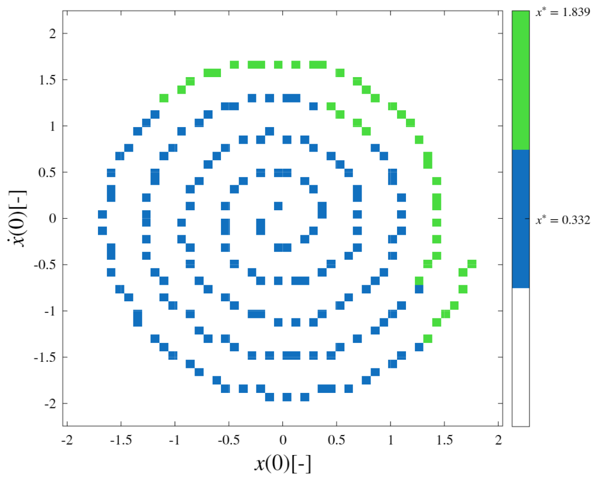

(a) Experimental spiral BoA.

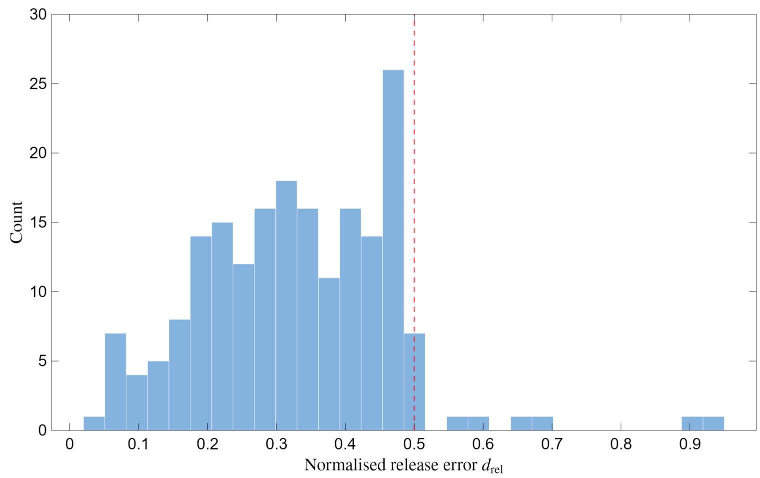

(b) Normalised release error distribution.

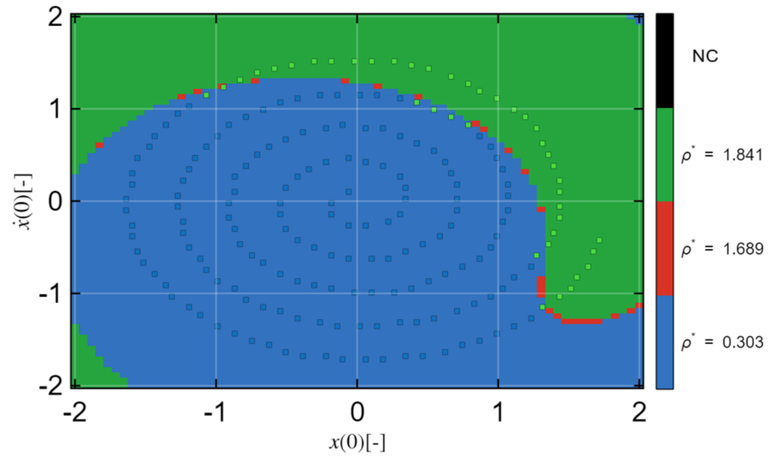

(c) Overlay on the corresponding numerical frequency-sweep frame of Figure 42.

Figure 59: Case 1, elliptic reference, higher forcing frequency. The erosion of the high-amplitude basin is reproduced experimentally.

5.3 Case 2: Strongly Nonlinear Forced Bistable System

Three stepping sequences are performed at the three frequencies selected in Section 4.4, namely near the lower fold, at mid-band, and near the upper fold, providing experimental counterparts to Figure 43. The results are shown in Figures 60, 61, and 62 respectively. In all three cases the normalised release error remains below the threshold, and the experimental basins overlay closely with the numerical references. The shifting dominance between the two coexisting attractors as the excitation frequency moves across the bistable band is correctly reproduced experimentally, confirming that the method tracks basin evolution without any retuning of the controller.

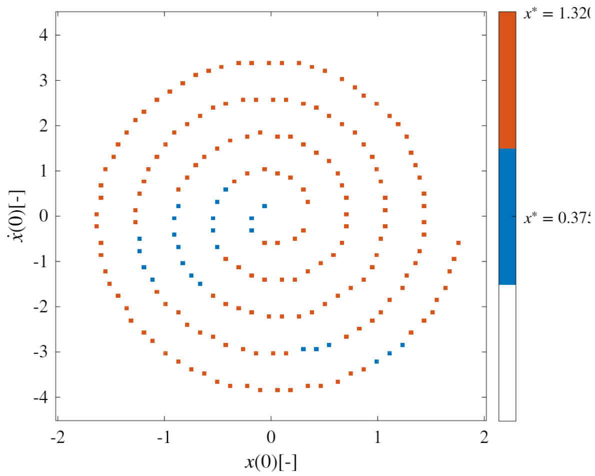

(a) Experimental spiral BoA (lower-fold frequency).

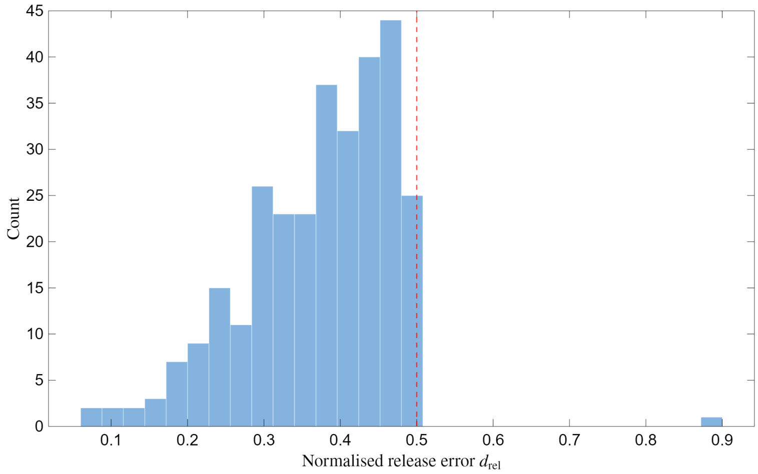

(b) Normalised release error distribution.

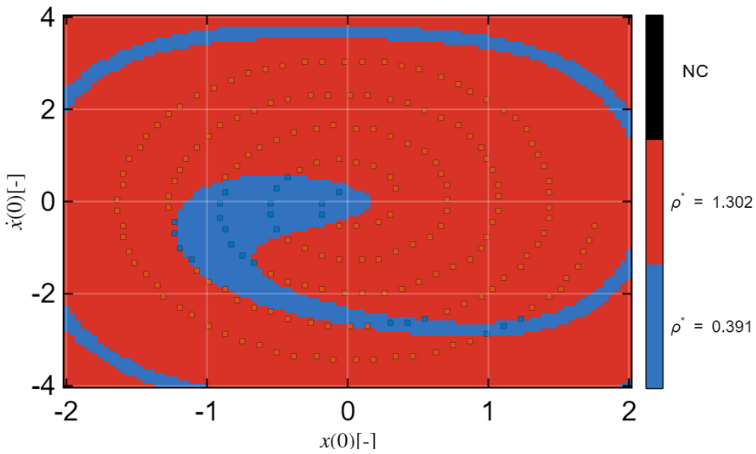

(c) Overlay on the corresponding numerical frame of Figure 43.

Figure 60: Case 2, elliptic reference, lower-fold frequency.

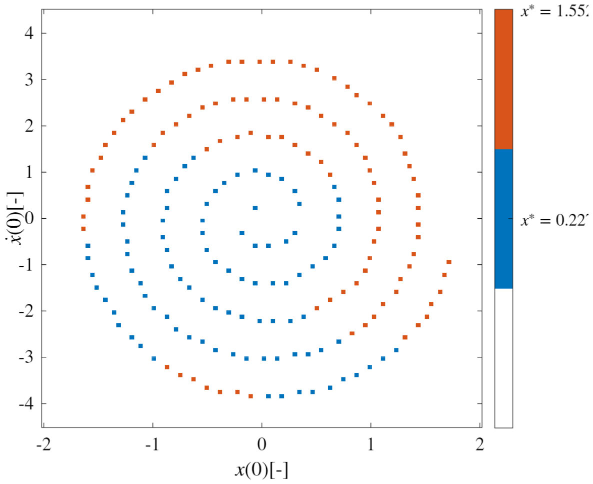

(a) Experimental spiral BoA ($\nu = 1.6$).

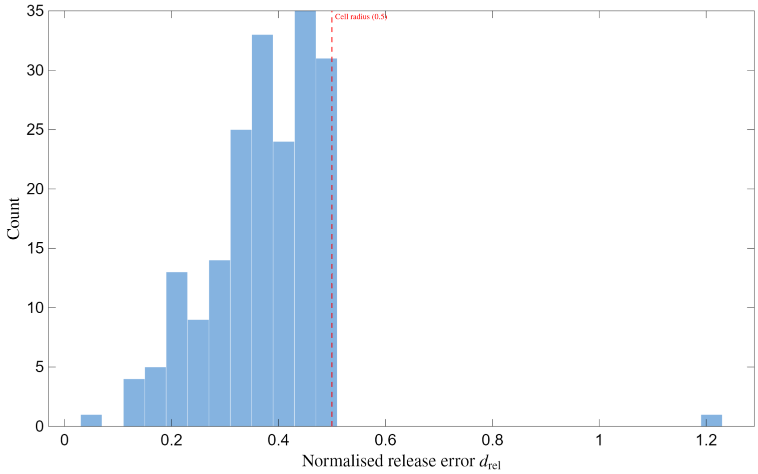

(b) Normalised release error distribution.

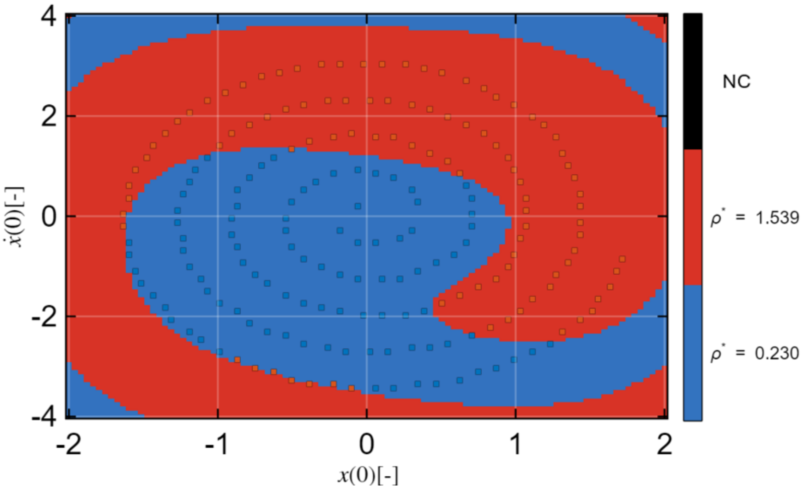

(c) Overlay on the numerical reference at $\nu = 1.6$.

Figure 61: Case 2, elliptic reference, nominal frequency $\nu = 1.6$.

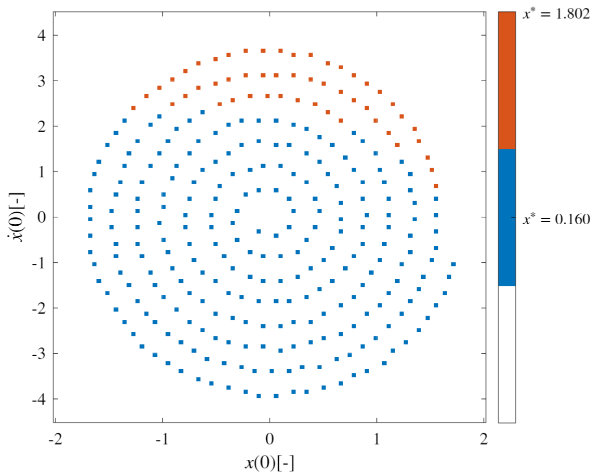

(a) Experimental spiral BoA (upper-fold frequency).

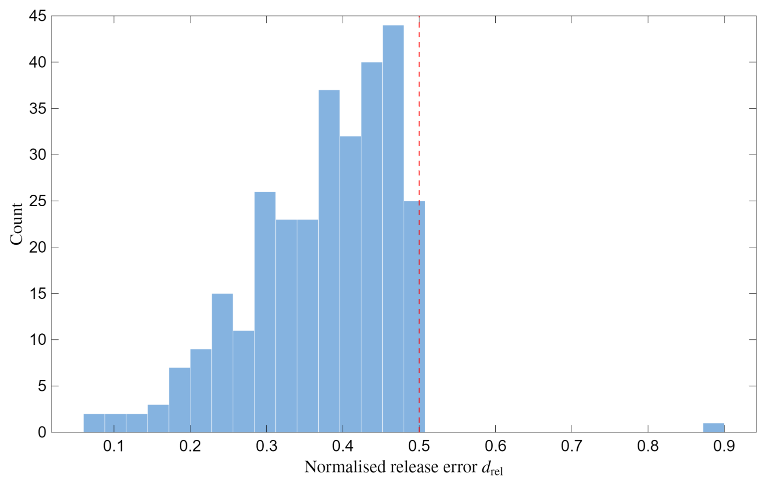

(b) Normalised release error distribution.

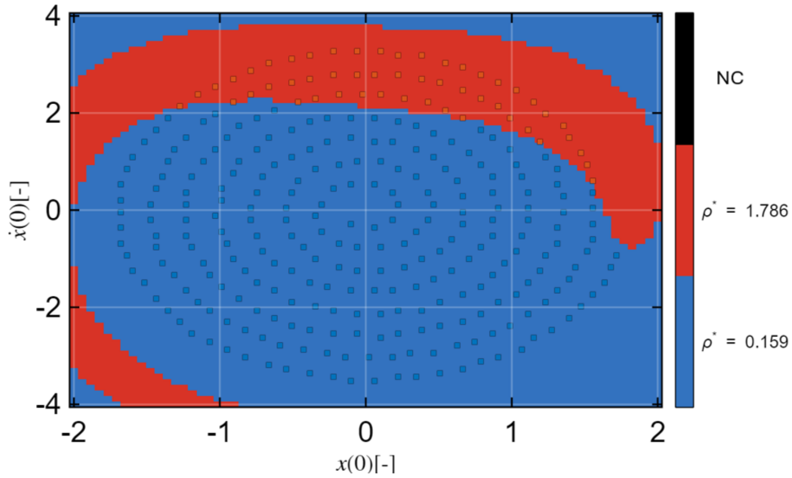

(c) Overlay on the corresponding numerical frame of Figure 43.

Figure 62: Case 2, elliptic reference, upper-fold frequency.

5.4 Case 3: Unforced Bistable System ($\alpha < 0$)

The linear stiffness potentiometer is tuned to its negative range and the DC-offset correction is applied before the stepping sequence begins. The dn reference is used for in-well cells and the cn reference for cross-well cells, following the switching strategy of Section 3.4.3.

The tracking controller successfully drives all initial conditions to their designated target cells. After release, the system settles into one of two stable equilibria measured at $x = 1.052$ and $x = -1.060$. This near-symmetrical distribution confirms that the DC offset has been correctly minimised. Figure 63 shows the experimental spiral BoA, the normalised release error distribution, and the overlay with the numerical reference.

When superimposing the experimental data onto the numerical basin, some deviations appear along the boundary between the two basins. The experimental stable equilibria deviate slightly from the ideal values of ± 1, indicating that the actual ratio of linear to cubic stiffness in the analogue hardware is slightly larger than nominal. This parametric shift alters the potential energy landscape and moves the curvature of the separating boundary. Furthermore, the unforced basin boundary is sensitive to the exact damping coefficient, since damping is the sole mechanism for energy dissipation in an autonomous system and governs how quickly a trajectory loses energy as it approaches an attractor. A small discrepancy between the physical damping and the nominal value changes this rate, so that trajectories started near the boundary can end up settling into the wrong well. It is strongly suspected that these localised mismatches are a direct consequence of physical parameter identification tolerances rather than a failure of the tracking algorithm.

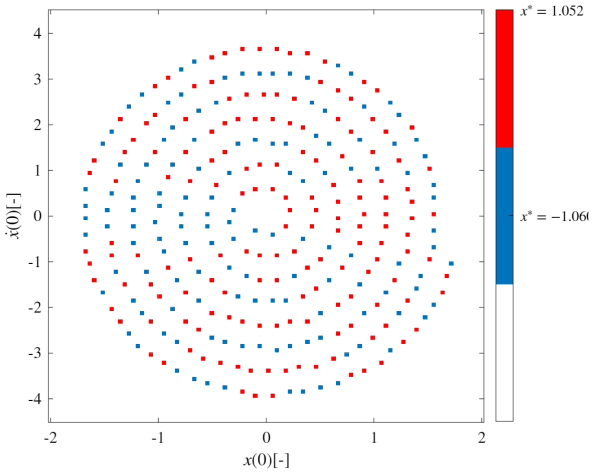

(a) Experimental spiral BoA.

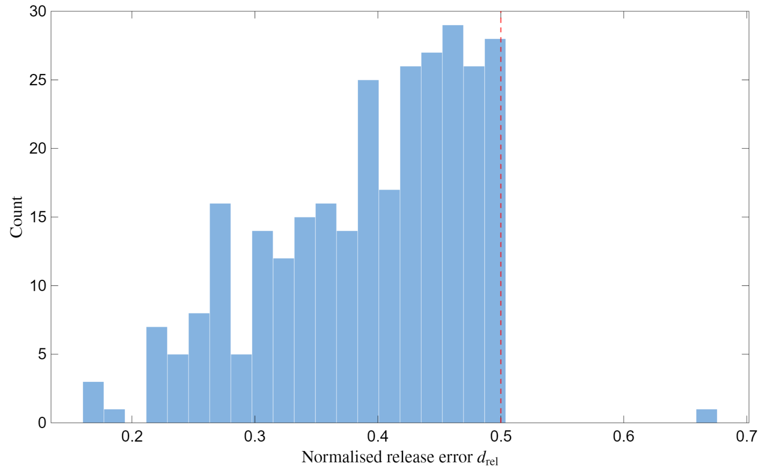

(b) Normalised release error distribution.

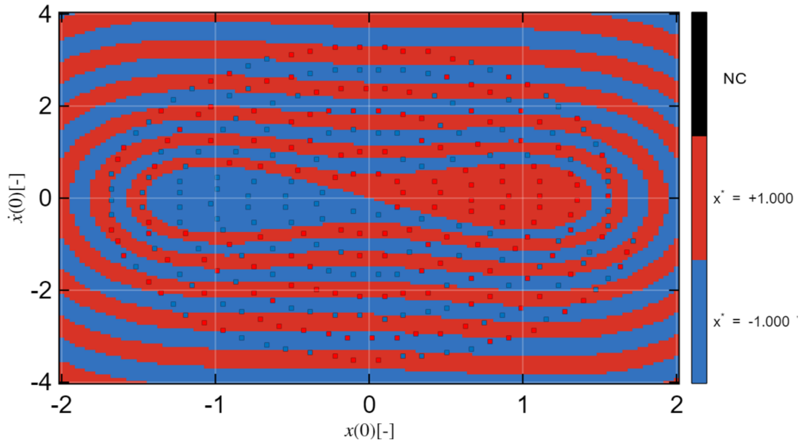

(c) Overlay on the numerical reference of Figure 39.

Figure 63: Case 3 experimental results (unforced bistable).

5.5 Case 4: Highly Forced System

5.5.1 Mid-Band Frequency

Case 4 presents a demanding validation scenario, primarily because of the coexisting subharmonic attractors whose basins must be correctly identified after release. The tracking controller functions correctly across all tested initial conditions, consistently driving the system to the targeted phase-space coordinates prior to release. After release, the steady-state detection algorithm successfully identifies the coexisting attractors predicted by the numerical model, including the subharmonic resonances. Figure 64 shows the experimental spiral BoA, the normalised release error distribution, and the overlay with the corresponding numerical frame of Figure 51. The overlay confirms high topological fidelity within both the primary domains and the localised subharmonic regions embedded in the low-amplitude basin.

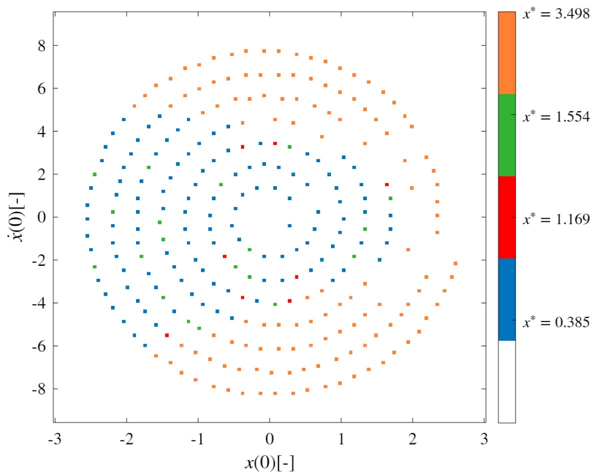

(a) Experimental spiral BoA (mid-band, subharmonics active).

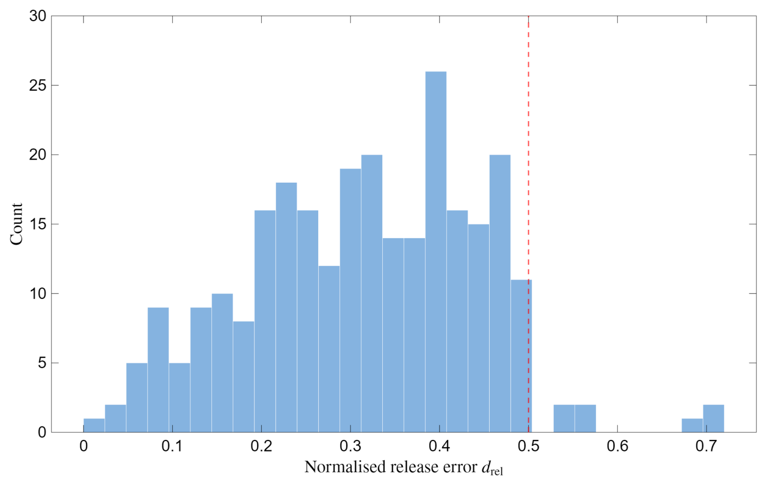

(b) Normalised release error distribution.

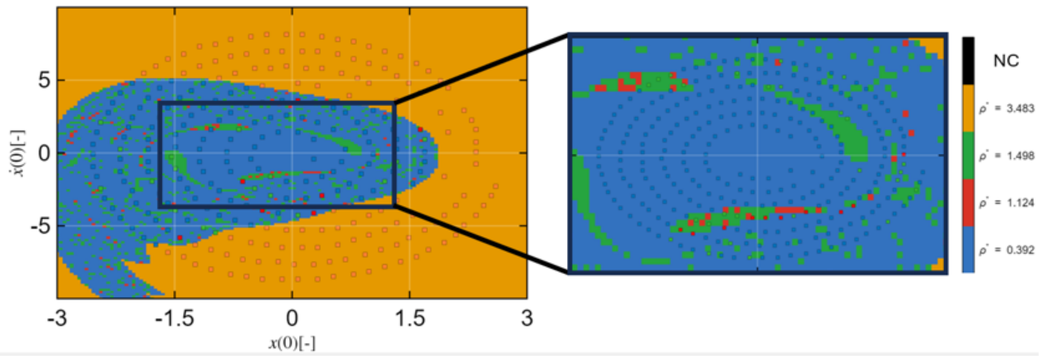

(c) Overlay on the corresponding numerical frame of Figure 51. The algorithm correctly distinguishes between the fundamental and subharmonic attractors.

Figure 64: Case 4, mid-band frequency (subharmonics active).

5.5.2 High-Band Frequency

The high-frequency stepping sequence tests the framework under a configuration where two subharmonic attractors coexist with the fundamental ones. The normalised release error remains below the threshold throughout, confirming accurate cell entry. After release, the detection routine successfully identifies the predicted coexisting $1:2$ and $1:3$ subharmonic attractors. The overlay shown in Figure 65 demonstrates close alignment with the numerical basin reference.

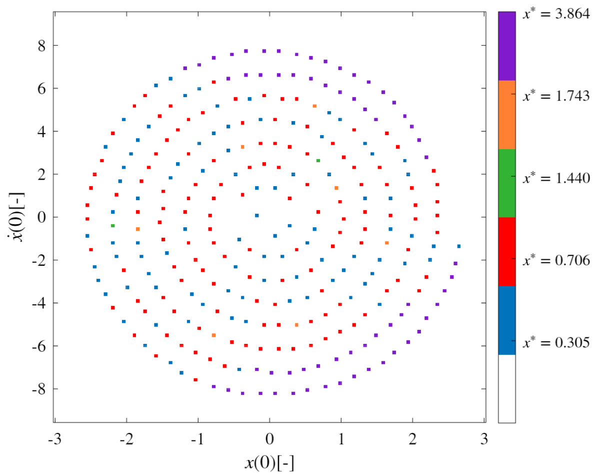

(a) Experimental spiral BoA.

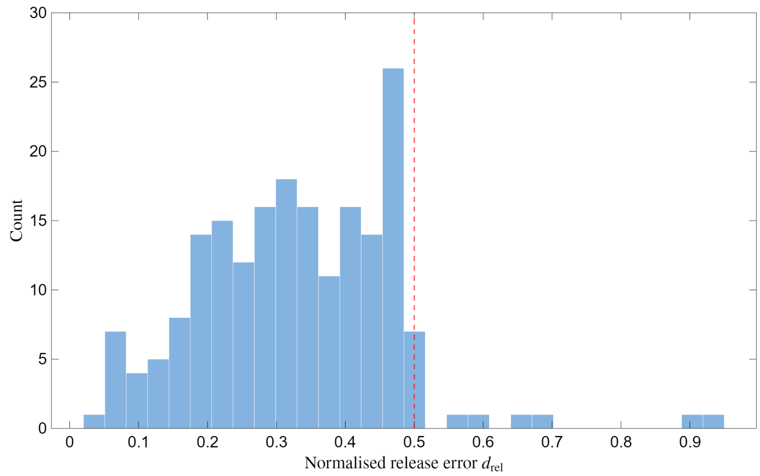

(b) Normalised release error distribution.

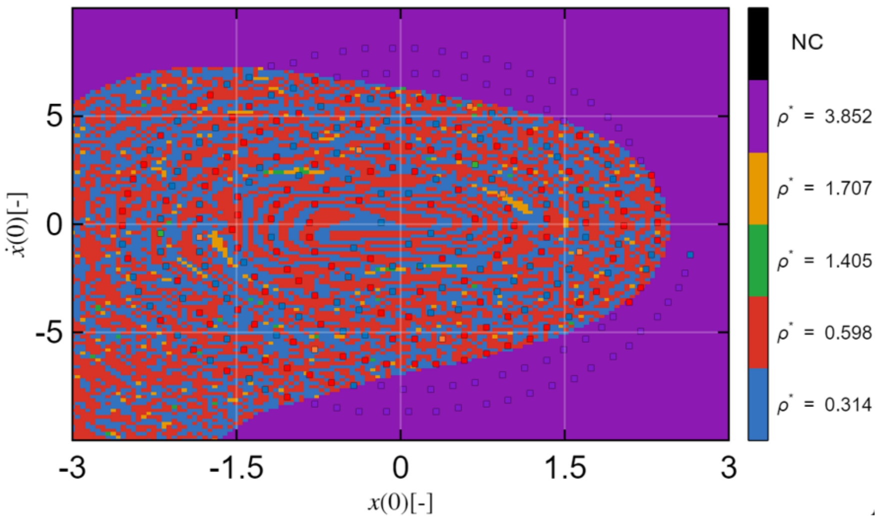

(c) Overlay on the corresponding numerical frame of Figure 51.

Figure 65: Case 4, high-band frequency.

5.6 Experimental Backfilling Results

The trajectory-based backfilling strategy described in Section 4.7 is applied experimentally to Cases 1, 2, and 3. On a 100×100 phase-space grid, executing full experimental runs for a subset of 400 initial conditions requires approximately 20 minutes. The backfilling strategy amplifies the data yield from this fixed experimental budget by propagating attractor labels to all grid cells whose centres are crossed by a confirmed stroboscopic path, without requiring any additional experimental runs for those cells.

For each case, Figure 66, 67, and 68 show the reconstructed basin, the classification status map, and the overlay with the numerical reference. In the status maps, cells resolved by direct experimental integration are shown in one colour, and cells resolved by backfilling are shown in another, making the efficiency gain immediately visible.

5.6.1 Case 1: Weakly Nonlinear Forced Bistable System

By running and backfilling 400 experimental initial conditions, the algorithm assigns attractor labels to approximately 40% of the entire grid, a tenfold increase over the 4% coverage that the same 400 cells would provide without label propagation. The overlay with the numerical reference confirms that the backfilled labels are correct throughout, with discrepancies confined to the basin boundary as in the spiral stepping results of Section 5.2.

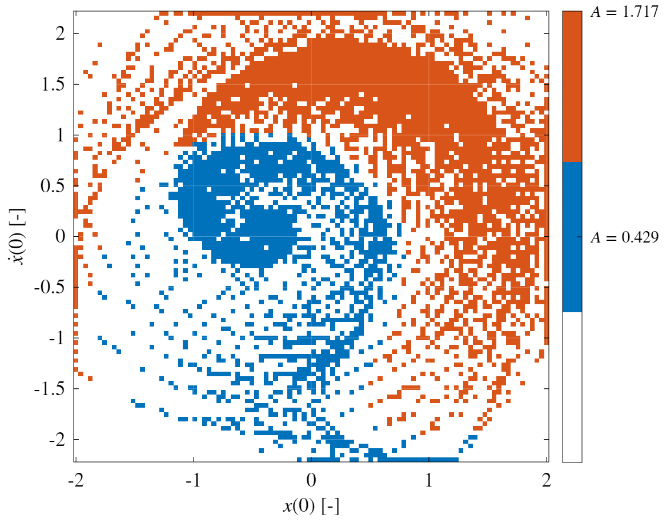

(a) Experimental backfilled BoA for Case 1.

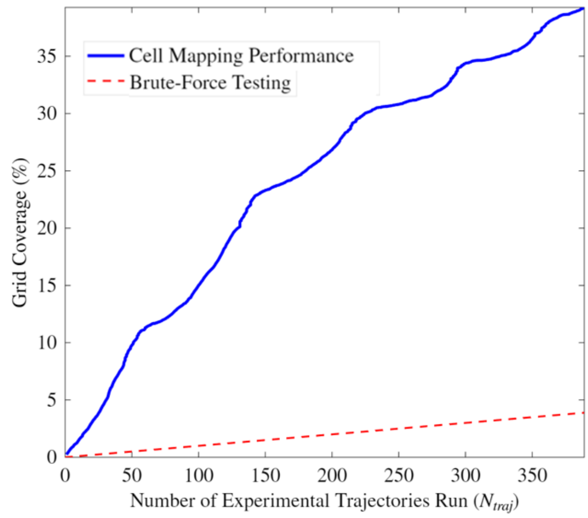

(b) Classification status: backfilling versus brute-force testing.

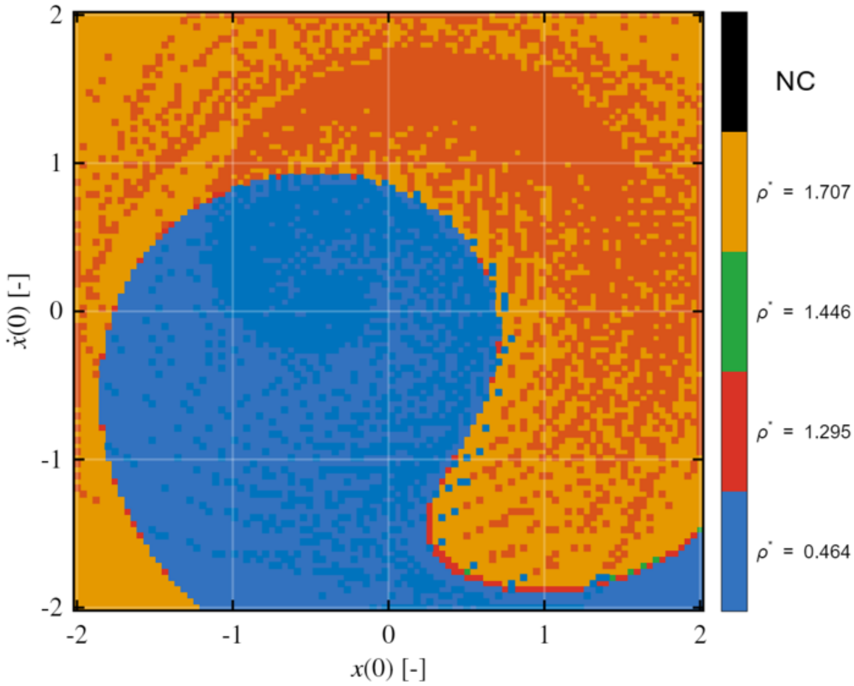

(c) Overlay on the numerical reference (Figure 29).

Figure 66: Case 1 experimental backfilling result ($\alpha = 1$, $\beta = 0.052$, $\delta = 0.016$, $\gamma = 0.039$, $\nu = 1.046$).

5.6.2 Case 2: Strongly Nonlinear Forced Bistable System

For Case 2, backfilling 400 experimental initial conditions yields attractor labels for approximately 12% of the grid, a threefold increase over the 4% baseline. The lower yield compared to Case 1 is a direct consequence of the higher damping coefficient ($\delta = 0.1$, nearly six times larger than in Case 1). Higher damping causes trajectories to settle faster, so each post-release path traverses fewer grid cells and the propagation reach of the backfill is shorter. Despite the reduced coverage, the overlay with the numerical reference shown in Figure 67 confirms close agreement.

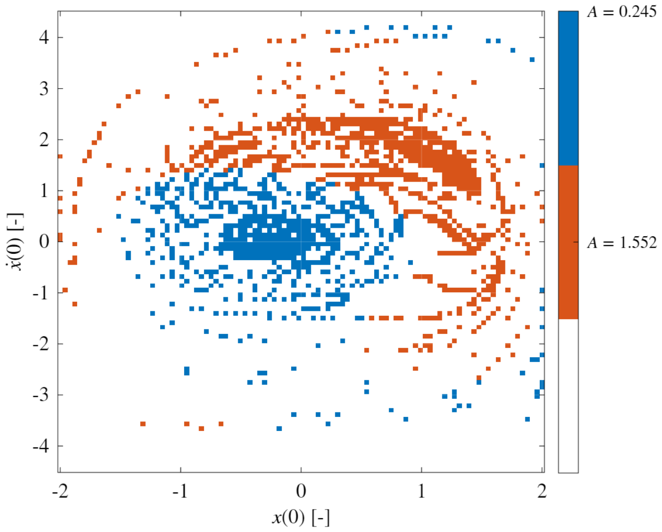

(a) Experimental backfilled BoA for Case 2.

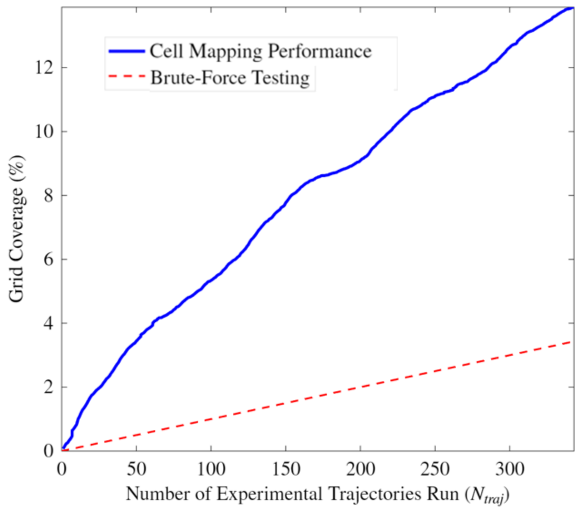

(b) Classification status: backfilling versus brute-force testing.

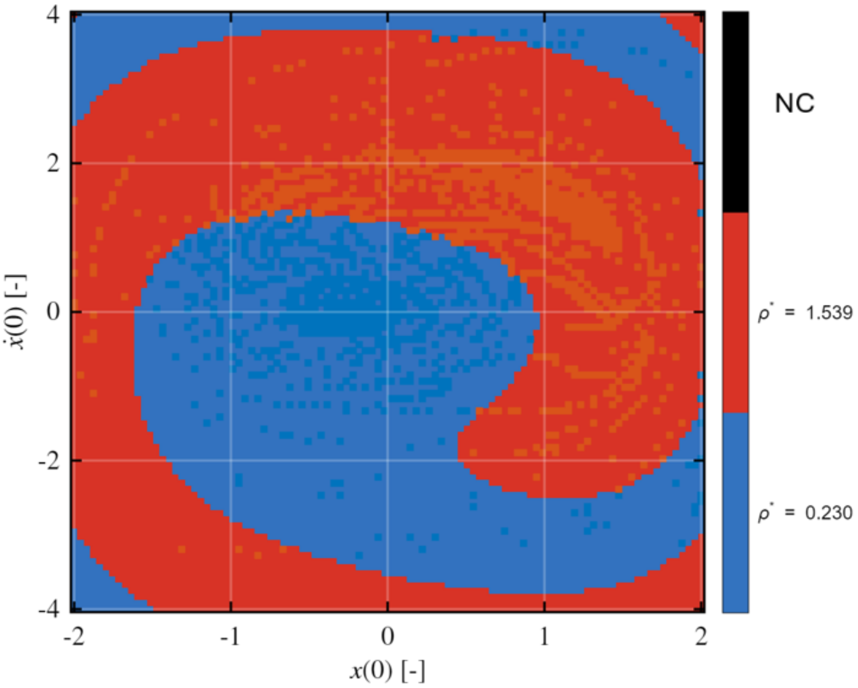

(c) Overlay on the numerical reference (Figure 38a).

Figure 67: Case 2 experimental backfilling result ($\alpha = 1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0.35$, $\nu = 1.6$).

5.6.3 Case 3: Unforced Bistable System ($\alpha < 0$)

For Case 3, 800 experimental runs with backfilling cover approximately 30% of the grid. Because the system is autonomous, all trajectories eventually settle into one of the two fixed-point attractors, and the long transients characteristic of the double-well system allow each trajectory to traverse a large number of grid cells before settling, which makes the backfill particularly effective here. The overlay shown in Figure 68 highlights the same parametric sensitivity discussed in Section 5.4: the experimental boundary deviates from the nominal numerical separatrix, particularly at larger initial conditions, confirming that small hardware parameter tolerances have a visible effect on the basin geometry in this autonomous configuration.

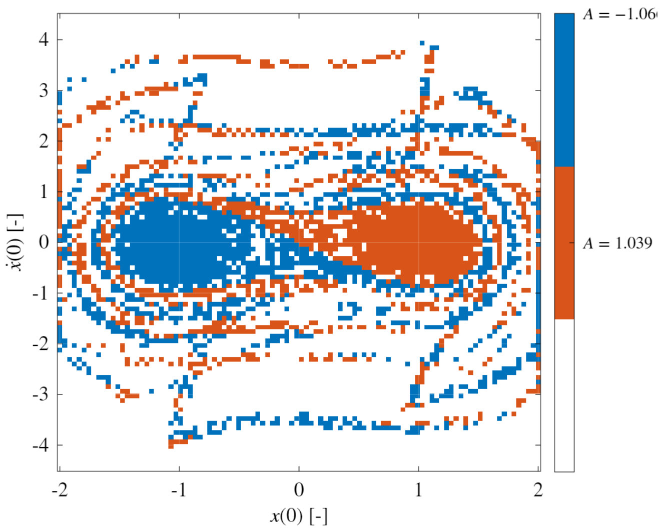

(a) Experimental backfilled BoA for Case 3.

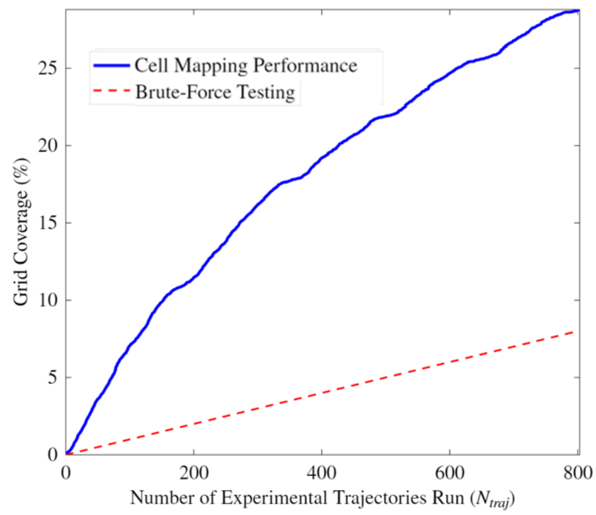

(b) Classification status: backfilling versus brute-force testing.

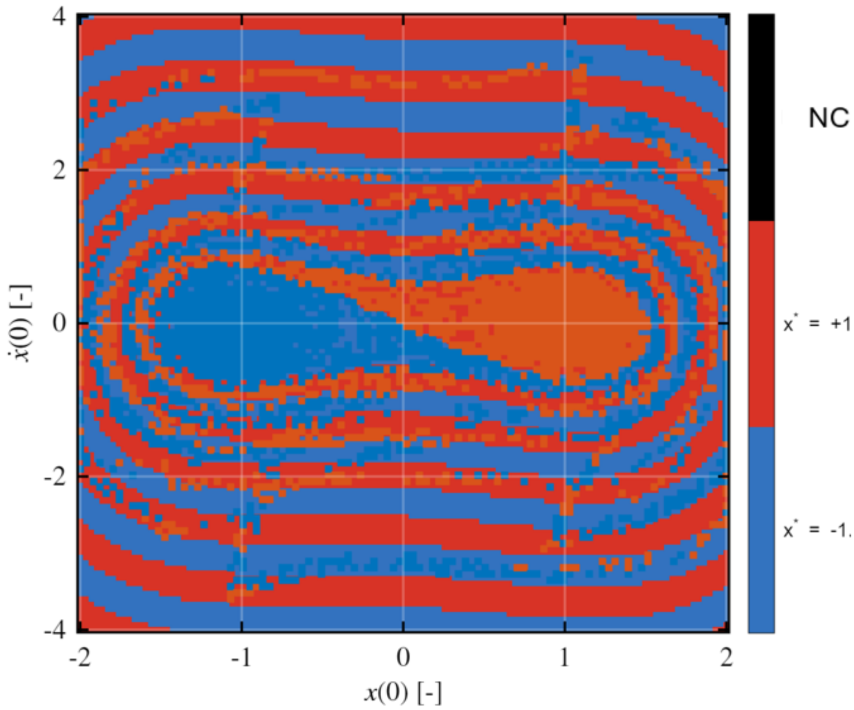

(c) Overlay on the numerical reference (Figure 39).

Figure 68: Case 3 experimental backfilling result ($\alpha = -1$, $\beta = 1$, $\delta = 0.1$, $\gamma = 0$).

5.7 Summary of Experimental Findings

The experimental campaign across Cases 1 to 4 demonstrates that the two-stage controlled method can be successfully executed on a physical dynamical system. The tracking controller consistently guided the circuit to the correct phase-space neighbourhood before release across all configurations, as confirmed by the normalised release error remaining below its threshold in every stepping sequence. After release, the steady-state detection correctly identified the attractor reached in each case, including configurations with coexisting attractors at the same forcing frequency.

The results across all four cases agreed well with the numerical references, with discrepancies confined to the basin boundary where small parameter tolerances have the most influence. The backfilling strategy provided a threefold to tenfold increase in grid coverage depending on system damping, without any additional experimental runs. Minor deviations observed in Case 3 along the separatrix are strongly suspected to stem from small differences between the physical hardware parameters and their nominal values rather than from any limitation of the algorithm itself, and they illustrate how sensitive the unforced double-well topology is to parameter accuracy.

6 Conclusion

6.1 Summary of Results

Basins of attraction describe which attractor a nonlinear system reaches depending on where it starts. Knowing this is important in practice: it tells the engineer which response will actually be observed, how sensitive that response is to disturbances, and whether a dangerous large-amplitude motion can be triggered accidentally from a nominally safe operating point. Despite this importance, measuring a basin experimentally has remained difficult, because it requires placing a physical system at many different, precisely specified states in the phase plane in a repeatable and automated way. This thesis proposed and demonstrated a method that addresses this problem.

The idea is to use feedback control to bring the system to each target state one by one, release it, and then observe which attractor it settles into. This two-stage procedure is repeated over a dense grid of initial conditions, building up the full basin map cell by cell. A key part of the method is the design of the reference trajectory that the controller tracks during the approach phase. A simple sinusoidal reference works well for low-amplitude targets, and a modulus criterion is used to decide when it is adequate. For larger amplitudes, references based on the exact conservative solutions of the Duffing equation are used instead, which ensures the reference stays on the correct energy level of the target and makes the tracking much more reliable. Control gains are computed automatically for each cell from the local dynamics, so no manual tuning is needed. A stroboscopic Poincaré criterion then monitors the free evolution after release and declares convergence once the system has settled.

The method was tested numerically on four Duffing configurations of increasing difficulty: a weakly nonlinear forced system, a strongly nonlinear forced system, an unforced double-well system, and a highly forced system exhibiting coexisting subharmonic attractors and disconnected response branches. In all four cases the reconstructed basins matched the reference maps closely, with differences confined to a narrow strip along the basin boundary. A series of tests at different frequencies confirmed that the basin evolution across the bistable band is tracked correctly without retuning the controller, and a sensitivity study showed that the method remains reliable under small parameter uncertainties. A backfilling strategy, which reuses the stroboscopic states visited along each trajectory to label additional grid cells, was shown to reduce the number of active integrations to 30 percent of the grid while keeping the result accurate.

The algorithm was then implemented in real time on a dSPACE MicroLabBox connected to an electronic analogue Duffing oscillator circuit, with no human intervention during the experimental sweeps. All four configurations were reproduced experimentally and the results agreed well with the numerical reference maps across all cases, which is the most important outcome of this work. The controlled initial-condition placement consistently guided the system to the correct phase-space neighbourhood before release, and the steady-state detection correctly identified the attractor reached in each case, including in configurations with coexisting attractors that a standard frequency sweep would not distinguish.

6.2 Original Contributions

The original contributions of this work are summarized as follows:

- **Automated Two-Stage Control Framework:** Development of a two-stage controlled method for the automated experimental reconstruction of basins of attraction on a Duffing oscillator. This framework combines precise simultaneous state placement (position and velocity) with systematic grid coverage, enabling the mapping of basin regions that are entirely inaccessible to standard uncontrolled methods.

- **Unified Reference Design Strategy:** Design of a versatile reference selection strategy for the initial stage that automatically switches between sinusoidal and Jacobi elliptic references based on target amplitude and dynamics. This covers monostable, cross-well bistable, and in-well bistable regimes within a single unified framework, using an elliptic modulus criterion as a practical decision rule.

- **Embedded Real-Time Implementation:** Successful real-time deployment of the complete algorithm on an embedded controller. This implementation includes an automated spiral sweep routine, an amplitude-scheduled gain table, and a Poincaré convergence detector, enabling fully unattended experimental mapping sessions on a physical analogue circuit.

- **Mapping of Isolated Responses:** Experimental detection and mapping of coexisting subharmonic attractors and hidden structures in the Duffing oscillator. This includes tracking completely disconnected response branches and isolated basin regions that remain completely invisible to conventional frequency sweeps and uncontrolled testing.

6.3 Directions for Future Research

The method relies on two pieces of knowledge that are not always available in practice. The first is the solution of governing equation in a conservative setting, which is used to build the Stage 1 reference trajectory. For the Duffing oscillator this solution is known analytically, but for most engineering nonlinear systems it is not. Without it, the reference would have to be computed numerically from an identified model, which introduces its own challenges. The second is the knowledge of the system parameters, which is needed both to design the reference and to schedule the control gains. In a real application these parameters must be identified beforehand, and any mismatch between the identified and true values will affect the accuracy of the initial-condition placement.

The method was also developed and demonstrated on a single-degree-of-freedom electronic circuit. Applying it to a physical mechanical structure would require adapting the actuation strategy for Stage 1 and accounting for the dynamics of the actuator in the gain design. For systems with more than one degree of freedom, sweeping a full grid becomes quickly impractical, and a targeted sampling strategy concentrating on the basin boundary would be needed.

Addressing these points, and in particular finding a way to design the reference without requiring the exact solution of the conservative system, is the most important direction for future work and would make the method applicable to a much wider class of systems.

References

[1] Nemanja Andonovski. Basins of attraction and dynamical integrity of nonlinear dynamical systems in high dimensions. PhD thesis, Università Politecnica delle Marche, 2019.

[2] Grégoire Bourdouch and Louise Marchant. FRC's codebase, 2025. GitLab repository, private academic access.

[3] Consortium des Equipements de Calcul Intensif (CECI). CECI – consortium des equipements de calcul intensif. `https://www.ceci-hpc.be`, 2025. Accessed: 2026-06-05.

[4] Thibaut Detroux, Jean-Philippe Noël, Laurent Masset, Gaëtan Kerschen, and Lawrence N. Virgin. Numerical study of intrinsic features of isolas in a 2-DOF nonlinear system. In Proceedings of the International Conference on Structural Engineering Dynamics (ICEDyn 2015), Lagos, Portugal, 2015.

[5] EEEGUIDE. Second order low pass Butterworth filter. `https://www.eeeguide.com/second-order-low-pass-butterworth-filter/`, 2016. Accessed: 2026-06-03.

[6] C. S. Hsu. A generalized theory of cell-to-cell mapping for nonlinear dynamical systems. Journal of Applied Mechanics, 48(3):634–642, 1981.

[7] 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. Special Issue: Non-linear Structural Dynamics.

[8] Ivana Kovacic and Michael J. Brennan, editors. The Duffing Equation: Nonlinear Oscillators and their Behaviour. John Wiley & Sons, Chichester, UK, 2011.

[9] Joseph P. LaSalle. The Stability of Dynamical Systems. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1976.

[10] Ali H. Nayfeh and Balakumar Balachandran. Applied Nonlinear Dynamics: Analytical, Computational, and Experimental Methods. Wiley, New York, 1995.

[11] Eddie Nijholt. The picard-lindelöf theorem: Existence and uniqueness of solutions, 2025. Lecture notes.

[12] Maarten Peeters, Régis Vigué, 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.

[13] Sergio P. Pellegrini, Nima Tolou, Mark Schenk, and Just L. Herder. Bistable vibration energy harvesters: A review. Journal of Intelligent Material Systems and Structures, 24(11):1303–1312, 2013.

[14] Ghislain Raze. An electronic Duffing oscillator. `https://github.com/GhislainRaze/Electronic-Duffing`, 2024. Accessed June 2026.

[15] Ghislain Raze, Gaëtan Abeloos, and Gaëtan Kerschen. Experimental continuation in nonlinear dynamics: recent advances and future challenges. Nonlinear Dynamics, 113:4949–4997, 2025.

[16] Ghislain Raze and Gaëtan Kerschen. A two-harmonic homotopy method (THHM) to connect a primary resonance to its secondary resonances. Proceedings of the Royal Society A, 481:20240875, 2025.

[17] M.S. Soliman and J.M.T. Thompson. Integrity measures quantifying the erosion of smooth and fractal basins of attraction. Journal of Sound and Vibration, 135(3):453–475, 1989.

[18] Alexandre Spits, Ghislain Raze, and Gaëtan Kerschen. A model- and derivative-free arclength continuation method for tracing bifurcation diagrams in nonlinear systems. In Journées annuelles du GdR EX-MODELI, Lille, France, 2025.

[19] J. C. Sprott and Anda Xiong. Classifying and quantifying basins of attraction. Chaos, 25(8):083101, 2015.

[20] Steven H. Strogatz. Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering. Studies in Nonlinearity. Westview Press, Boulder, CO, 2 edition, 2015.

[21] R. Wiebe, Lawrence N. Virgin, and S. M. Spottswood. Stochastic interrogation of competing responses in a nonlinear distributed system. Nonlinear Dynamics, 79(1):607–615, 2015.

[22] Wenzhong Yan, Yunchen Yu, and Ankur Mehta. Analytical modeling for rapid design of bistable buckled beams. Theoretical and Applied Mechanics Letters, 9(4):264–272, 2019.