

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Algorithmic Limits in Optimal Control: A Computational and Statistical Study

Permalink

<https://escholarship.org/uc/item/7x048410>

ISBN

9798186381952

Author

Rajabi, Faranak

Publication Date

2026-05-26

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Santa Barbara

**Algorithmic Limits in Optimal Control:
A Computational and Statistical Study**

A Thesis submitted in partial satisfaction
of the requirements for the degree

Master of Science

in

Computer Science

by

Faranak Rajabi

Committee in charge:

Professor Frédéric Gibou, Chair

Professor Xifeng Yan

Professor Francesco Bullo

June 2026

The thesis of Faranak Rajabi is approved.

Professor Xifeng Yan

Professor Francesco Bullo

Professor Frédéric Gibou, Committee Chair

June 2026

Algorithmic Limits in Optimal Control:
A Computational and Statistical Study

Copyright © 2026

by

Faranak Rajabi

In memory of Matin Khademi.

A young Iranian student in Turin, Italy, who ended his own life under the unbearable psychological pressure of the regime's massacre of the Iranian people. Iranians living in Turin held his funeral without his family present. Because of the internet blackout imposed inside Iran, his family could not even see pictures of their son's burial in another country.

The crimes of this regime will be forever recorded in history.

Javid Shah.

Acknowledgements

I want to thank my advisors, Prof. Jeff Moehlis and Prof. Frédéric Gibou, for their patient guidance through this unfamiliar path, and my committee members, Prof. Xifeng Yan and Prof. Francesco Bullo, for their time and feedback. I am also grateful to Prof. João Hespanha and Prof. Linda Petzold for what they taught me along the way.

I want to thank my parents, who have never been able to attend my defenses or graduations, and who will not have access to the internet to attend this one either. I know they are the ones who wanted to be here most. I thank my brother Mohammad and my sister Fatemeh, whose love has been an anchor across thousands of miles.

I want to thank my friends, both at home and here. Ghazaleh in Tehran, my friend of eleven years. Farid in London. Omid in Switzerland, whose insights into control theory shaped this thesis. And my community in Santa Barbara: Yeganeh, Parvin, Danial, Arghavan, Parsa, Ryan, Nimas(Both Alizadeh and Kelidari!), Alvis, Sepideh, Javad, and Hojjat, who was my first coding teacher. And Ponyo, my sweet little kitty, for the comfort she brings.

This thesis was written during a period of extraordinary hardship for the Iranian people. Since I began this work in September 2023, my country has witnessed two wars, a revolution, and a brutal massacre. I dedicate this work to the Iranian people who were killed by the Islamic regime during the recent unrest, to their families, and to those who survived the streets but witnessed the deaths of their friends and children. I dedicate it to the family of Matin Khademi, the Iranian master's student in Italy who took his own life under the pressure of these events, and whose parents could not attend his funeral. I hope his mother finds peace. My compatriots and I will never be the same after what has happened to our nation. We will not forget, and we will not forgive.

Curriculum Vitæ

Faranak Rajabi

Education

- 2026** M.S. in Computer Science, University of California, Santa Barbara, USA.
- 2027** Ph.D. in Mechanical Engineering (Expected), University of California, Santa Barbara, USA.
- 2021** B.S., Sharif University of Technology, Tehran, Iran.

Publications

Journal papers:

1. F. Rajabi, J. Fingerman, A. Wang, J. Moehlis, F. Gibou. CASL-HJX: A comprehensive guide to solving deterministic and stochastic Hamilton-Jacobi equations, *Computer Physics Communications*, vol. 319, 109941, 2026.
2. F. Rajabi, F. Gibou, J. Moehlis. Optimal control for stochastic neural oscillators, *Biological Cybernetics*, vol. 119, Article 9, 2025.
3. J. Moehlis, M. Zimet, F. Rajabi. Magnitude-constrained optimal chaotic desynchronization of neural populations, *Frontiers in Network Physiology*, accepted 2025.
4. F. Rajabi, F. Gibou. Solving advection-diffusion-reaction equations with Robin boundary conditions on piecewise smooth interfaces, *Journal of Computational Physics*, 2026 (under review).

Conference papers:

5. M. Zimet, F. Rajabi, J. Moehlis. Nearly optimal chaotic desynchronization of neural oscillators, *Proceedings of the 64th IEEE Conference on Decision and Control (CDC)*, 2025.
6. F. Rajabi, F. Gibou, J. Moehlis. Beyond phase reduction: Amplitude collapse in optimal control of coupled oscillators, *65th IEEE Conference on Decision and Control (CDC)*, 2026 (under review).

Abstract

Algorithmic Limits in Optimal Control: A Computational and Statistical Study

by

Faranak Rajabi

Optimal control of nonlinear dynamical systems is often approached through reduced-order models, gradient-based optimization, or stochastic search methods. While computationally tractable, these approaches implicitly restrict where optimal solutions may reside in state space. This thesis studies coupled nonlinear oscillators and shows that these assumptions can fail simultaneously. By solving the full four-dimensional Hamilton-Jacobi-Bellman equation on a GPU-accelerated grid, we identify a critical coupling strength beyond which the globally optimal control leaves the limit-cycle manifold entirely. This regime is missed by phase reduction, gradient methods, and stochastic search alike, establishing the transition as a multi-method failure boundary rather than a limitation of any one approach.

The thesis makes two main contributions. First, it establishes off-manifold optimal control as a fundamentally full-state phenomenon and identifies a mechanism, termed *amplitude collapse*, through which optimal trajectories temporarily suppress oscillator amplitudes to reduce coupling-induced constraints. Analytical predictions based on a phase-velocity surplus argument agree closely with numerical results and generalize beyond Stuart-Landau oscillators to coupled Van der Pol systems. Second, the thesis develops a GPU-accelerated four-dimensional Hamilton-Jacobi solver extending the CASL-HJX framework with fifth-order WENO reconstruction, SSP-RK3 time integration, and CUDA-based Hamiltonian evaluation. The implementation computes full-state value functions at 64^4 resolution in approximately five minutes on a single NVIDIA A100 GPU.

Together, these results demonstrate that optimizing within a prescribed reduced model can systematically miss the true global optimum and that modern GPU-accelerated Hamilton-Jacobi methods now make full-state optimal control computation feasible for nontrivial coupled nonlinear systems.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Thesis Contributions	2
1.3	Outline	4
1.4	Permissions and Attributions	5
2	Coupled Oscillators and the Antiphase Locking Problem: Background and Literature	6
2.1	Coupled Limit-Cycle Oscillators in Science and Engineering	6
2.2	The Antiphase Locking Problem	7
2.3	Phase Reduction and Its Extensions	8
2.4	Stochastic Oscillators	9
2.5	The Gap Addressed by This Thesis	10
3	Computational Methods for Oscillator Optimal Control: A Comparative Survey	11
3.1	Direct Optimization Methods	11
3.2	Reduced-Order Methods	13
3.3	Full-State Methods	14
3.4	Method Selection: The Problem This Thesis Addresses	16
4	Why Solve the Hamilton-Jacobi-Bellman Equation in Four Dimensions	17
4.1	Statement of the Problem	18
4.2	Why Not Phase Reduction Alone	20
4.3	Why Not Direct Optimization in $\mathbb{R}^4$	21
4.4	Two Approaches in Comparison	22
4.5	The Hamilton-Jacobi-Bellman Equation	23
4.6	Cost: Four Dimensions Is the Boundary of Feasibility	25
5	Numerical Methods for the Hamilton-Jacobi-Bellman Equation	26
5.1	The Hamiltonian and Optimal Feedback	26

5.2	Spatial Discretization: Local Lax-Friedrichs WENO5	27
5.3	Time Integration: SSP-RK3	29
5.4	Solver Pseudocode	29
5.5	Boundary Conditions	30
5.6	Verification: What Each Test Proves	30
6	GPU Acceleration and the Extension of CASL-HJX to Four Dimensions	33
6.1	Computational Cost of 4D HJB	33
6.2	Parallelization Strategy	34
6.3	Extension from 2D to 4D	35
6.4	Performance Profile	35
6.5	Code Availability and Software Architecture	36
7	Deterministic Results: Amplitude Collapse and Multi-Method Failure	39
7.1	Problem Formulation	40
7.2	HJB Solution Method	42
7.3	Two Qualitatively Distinct Strategies	43
7.4	Sharp Strategy Transition	45
7.5	Physical Mechanism: Phase Velocity Surplus	46
7.6	Role of Nonisochronicity	47
7.7	Quantitative Verification: Free Phase Budget	48
7.8	Comparison with Direct Optimization	48
7.9	Comparison with Phase-Reduced Control	49
7.10	Validation on Coupled Van der Pol Oscillators	51
7.11	Summary	52
8	Conclusions and Future Directions	57
8.1	Summary of Contributions	57
8.2	Implications	58
8.3	Limitations	59
8.4	Future Directions	60
8.5	Closing Remarks	61
	Bibliography	62

Chapter 1

Introduction

1.1 Motivation

Coupled limit-cycle oscillators arise across science and engineering, and controlling their phase relationships has direct consequences in clinical neurostimulation and power-grid infrastructure. In deep brain stimulation for Parkinson’s disease, pathological synchrony of neural populations drives motor symptoms, and current open-loop protocols (high-frequency pulse trains above 130 Hz) consume sufficient energy that implanted batteries require surgical replacement every three to five years [1, 2]. In power systems, grid-forming inverters increasingly employ virtual oscillator control based on coupled Hopf-normal-form dynamics [3, 4], and desynchronization in this setting can cascade to system-wide failure [5]. In both domains, the underlying control problem is the same: steer coupled oscillators toward a desired phase relationship while minimizing actuation energy.

The standard tool for this problem is phase reduction [6, 7], which projects each oscillator onto a single phase variable evolving on its limit cycle. Phase reduction has enabled elegant optimal control formulations and closed-form solutions in the weak-input regime [8, 9], and recent work has extended the framework using phase-amplitude and

isostable coordinates [10, 11, 12, 13]. By construction, however, these methods constrain the state to lie on or near the limit cycle. When coupling or forcing becomes strong, pure phase descriptions break down and the optimal control mechanism itself may lie outside the manifold underlying these methods.

This thesis examines that boundary. Working with two diffusively coupled Stuart-Landau-like oscillators as a tractable but representative case study, we solve the full four-dimensional Hamilton-Jacobi-Bellman equation, characterize where reduced-order methods and direct optimization fail, and develop statistical tools for certifying those failures in advance.

1.2 Thesis Contributions

The thesis makes the following contributions:

1. **GPU-accelerated 4D HJB solver.** We extend the CASL-HJX framework [14] to four spatial dimensions with GPU acceleration, implementing 5th-order WENO spatial reconstruction, SSP-RK3 time integration, and CUDA kernels templated on problem-specific Hamiltonians. The implementation enables full-state HJB solutions on a $G = 64^4$ grid in approximately five minutes on a single NVIDIA A100, against an estimated ten to twenty hours on a comparable CPU implementation. The solver is validated through (i) 5th-order convergence on linear advection in L^2 and L^∞ , the test that genuinely probes the WENO5 order because the exact solution is smooth and the order appears in the error decay rate; (ii) agreement with the analytical algebraic Riccati equation solution for a 4D double integrator (relative L^2 error of 6×10^{-5} at $G = 64$), which serves as an end-to-end pipeline check rather than an order check, since the LQR value function is quadratic in the state and is captured

exactly by any sufficiently high-order scheme; and (iii) grid convergence on the Stuart-Landau problem itself.

2. **Off-manifold optimal control mechanism.** Using the solver, we identify a critical coupling $c^* = 0.227 \pm 0.003$ above which the globally optimal control mechanism for the coupled oscillator problem transiently collapses one oscillator's amplitude toward the origin, neutralizing coupling opposition and exploiting nonisochronous phase drift before rebuilding amplitude. The mechanism is verified analytically through a phase-velocity-surplus argument, $\delta\omega = b^2/(\omega + \Omega)$, and quantitatively through the computed value function and feedback law, with the predicted free-phase budget matching the measured value to within 1° across a sweep over nonisochronicity.
3. **Multi-method failure characterization.** We show that this off-manifold mechanism is not recovered by gradient-based optimization (L-BFGS-B), by global stochastic search (CMA-ES with up to 1.96×10^5 function evaluations), or by phase-reduced optimal control with or without an interaction function. The same critical coupling c^* marks the failure boundary for all three families of methods, establishing the transition as a structural feature of the problem rather than an algorithmic artifact.
4. **Falsification on a non-Stuart-Landau system.** We verify that the off-cycle mechanism is not specific to the Stuart-Landau algebraic structure by solving the same control problem on coupled Van der Pol oscillators. The off-cycle excursion persists (Property 1: diffusive coupling vanishes at the origin in any limit-cycle system) while the cost non-monotonicity disappears (Property 2: phase-velocity surplus requires nonisochronicity, which Van der Pol lacks). This two-by-two factorial cleanly separates the geometric prerequisite for off-cycle control from the algebraic prerequisite for the associated energetic advantage.

The four contributions correspond to completed work that has been submitted to the 65th IEEE Conference on Decision and Control (CDC 2026).

1.3 Outline

Chapter 2 reviews the literature on coupled oscillator control, phase reduction, and the antiphase locking problem, and identifies the gap addressed by this thesis. Chapter 3 surveys computational methods that have been applied to oscillator optimal control, with discussion of their respective strengths and limitations. Chapter 4 motivates the choice of full-state HJB over reduced-order alternatives for the strong-coupling regime, and develops the HJB equation, its Hamiltonian decomposition, and the closed-form saturated feedback law in full detail. Chapter 5 presents the numerical methods underlying our HJB solver, including the WENO5 spatial discretization, SSP-RK3 time stepping, explicit solver pseudocode for one backward step, and verification against three test problems with a careful discussion of what each test actually proves about the solver. Chapter 6 describes the GPU acceleration strategy and the extension of CASL-HJX from two to four spatial dimensions. Chapter 7 presents the deterministic results, including the sharp strategy transition at $c^* = 0.227$, the phase-velocity-surplus mechanism, the saturating-optimal structure of the resulting control, the comparison with direct optimization and phase-reduced methods, and the validation on coupled Van der Pol oscillators (this content corresponds to the work submitted to CDC 2026). Chapter 8 concludes and outlines directions for future work.

1.4 Permissions and Attributions

1. The content of Chapter 7 is the result of a collaboration with F. Gibou and J. Moehlis, and has been submitted for publication at the IEEE Conference on Decision and Control (CDC) 2026 under the title “Beyond Phase Reduction: Amplitude Collapse in Optimal Control of Coupled Oscillators” [15].
2. The numerical infrastructure used throughout this thesis builds on the CASL-HJX framework, which is the result of a collaboration with J. Fingerman, A. Wang, J. Moehlis, and F. Gibou, and has been published in Computer Physics Communications [14].

Chapter 2

Coupled Oscillators and the Antiphase Locking Problem: Background and Literature

2.1 Coupled Limit-Cycle Oscillators in Science and Engineering

Limit-cycle oscillators provide a canonical mathematical model for self-sustained periodic activity in systems ranging from neuronal spiking to mechanical resonance to electrochemical reactions. When such oscillators are coupled, their phase relationships become objects of central interest because synchronization and its breakdown drive collective behavior. Three application domains motivate the present study.

Neural synchronization and Parkinson's disease. Pathological synchronization of neuronal activity in the basal ganglia, particularly within the beta frequency band, is a hallmark of motor symptoms in Parkinson's disease [16, 17, 18]. Deep brain stimulation

(DBS) at the subthalamic nucleus or globus pallidus interna provides therapeutic benefit by disrupting this synchrony, but the conventional high-frequency open-loop protocol consumes substantial energy and requires periodic surgical battery replacement [1, 19]. Closed-loop and adaptive DBS approaches seek to reduce energy consumption by responding to physiological biomarkers [16, 20, 21], and optimal control formulations have been proposed in both deterministic and stochastic settings [22, 23, 24].

Power-grid inverter control. Modern power systems are increasingly dominated by inverter-based resources whose synchronization with the grid is a critical stability concern. Virtual oscillator control schemes based on coupled Hopf-normal-form dynamics provide grid-forming behavior with desirable robustness properties [3, 4]. Loss of synchronization in such systems can cascade to network-wide failure, and the design of energy-efficient stabilizing controllers remains an active research area [5].

Coupled electrochemical and laser systems. Coupled electrochemical oscillators [9] and coupled-laser networks have served as experimental platforms for testing phase-control theory. These systems offer the advantage of tunable coupling and clean diagnostics, and have been used to validate weak-coupling phase models.

2.2 The Antiphase Locking Problem

The antiphase state, in which two oscillators maintain a phase difference of π radians, is a canonical target for desynchronization control. It is the natural “opposite” of in-phase synchronization, and for many coupled oscillator systems it is a stable solution under the natural dynamics for appropriate parameters [25, 6]. Driving an initially in-phase pair to the antiphase configuration while minimizing actuation energy is therefore a well-posed optimal control problem with both theoretical and practical significance [8, 22, 26].

For coupled oscillators with diffusive coupling and nonisochronous phase dynamics,

the antiphase configuration has a finite basin of attraction. Steering the system into this basin requires precise control, and the small basin size makes the problem sensitive to both modeling assumptions and computational method. This sensitivity is one reason the antiphase problem provides a useful diagnostic for evaluating optimal control methodology.

2.3 Phase Reduction and Its Extensions

Phase reduction is the cornerstone of analytical and computational oscillator control [6, 27, 7]. The central idea is that a d -dimensional oscillator near a stable limit cycle can be described, to leading order in the perturbation, by a single phase variable ϕ on the cycle. The infinitesimal phase response curve (iPRC), defined as the gradient of the asymptotic phase with respect to the state, governs how external inputs modify the phase. For weakly coupled oscillators, the dynamics reduce further to a phase-difference equation governed by an interaction function $H(\psi)$ [6, 25].

This framework has supported several decades of work on oscillator control. Optimal phase-control problems can be cast as variational problems over the phase coordinate and solved by Pontryagin’s principle [8, 9]. Inputs can be designed for entrainment, phase locking, or desynchronization [28, 29].

The limitations of phase reduction have been well-documented. Stronger inputs invalidate the small-perturbation assumption and force a richer description [30, 11, 31]. Several lines of work have extended the framework to broader regimes.

Augmented phase reduction retains the phase variable but appends amplitude or isostable coordinates that capture transverse dynamics [26, 10]. Isostable coordinates are level sets of slowest-decaying eigenfunctions of the linearized return map and provide a principled extension beyond the cycle.

Adaptive phase-amplitude coordinates construct curvilinear coordinate systems

that remain valid for larger excursions from the cycle, and have been used to design optimal controls for strongly perturbed oscillators [12, 32, 13, 33].

Phase-isostable coordinates for coupled systems address the problem of network dynamics under stronger coupling [34, 35].

Despite these extensions, all of these methods retain the structural assumption that the optimal trajectory lives in a neighborhood of the cycle. The thesis demonstrates a regime in which the globally optimal trajectory exits this neighborhood entirely.

2.4 Stochastic Oscillators

When intrinsic or extrinsic noise is non-negligible, the deterministic phase reduction must be replaced by a stochastic counterpart. Several frameworks exist.

The classical approach treats noise as a small perturbation and computes phase diffusion coefficients via the Fokker-Planck equation [6]. More recent work by Pérez-Cervera and collaborators [36] provides a universal description of stochastic oscillators in terms of eigenfunctions of the Kolmogorov backward operator, defining a noise-respecting stochastic phase that remains well-defined globally, including at the deterministic phaseless set. This framework is particularly relevant to the present thesis because the off-manifold mechanism we identify operates precisely in regions where the deterministic phase is ill-defined.

Optimal control of stochastic oscillators has been studied via stochastic HJB methods [23, 24] and via path-integral and sample-based dynamic programming approaches. Extending these tools to higher-dimensional coupled systems is an active research direction beyond the scope of this thesis.

2.5 The Gap Addressed by This Thesis

The literature reviewed above establishes two facts. First, phase reduction and its extensions are powerful tools for oscillator control, but they assume that the optimal solution lies near the limit cycle. Second, when this assumption fails, the existing literature does not provide a constructive characterization of the resulting optimal mechanism for coupled diffusively-coupled oscillators, nor does it provide a calibrated diagnostic for when the assumption is reliable.

This thesis closes the first gap. We characterize an off-cycle optimal mechanism that no reduced-order method in the reviewed literature recovers.

Chapter 3

Computational Methods for Oscillator Optimal Control: A Comparative Survey

This chapter surveys the computational methods that have been applied to oscillator optimal control problems and discusses their respective strengths and limitations. The comparison establishes the context for the methodological choice made in Chapter 4 to use full-state HJB rather than a reduced-order or direct optimization approach.

3.1 Direct Optimization Methods

Direct optimization approaches parameterize the control input $u(t)$ and minimize the cost functional over the parameters using numerical optimization. The optimization problem is finite-dimensional in the control parameters, and the output is a single optimal trajectory associated with one initial condition. This class of methods is widely used because it is conceptually simple and broadly applicable.

3.1.1 Gradient-Based Methods

When the cost gradient with respect to the control parameters is available, either by adjoint sensitivity analysis or by automatic differentiation, gradient-based methods such as L-BFGS-B [37] are computationally efficient. They typically require $O(n)$ function evaluations per iteration, where n is the number of control parameters, and exhibit superlinear local convergence near a minimum.

Pros: efficient when gradients are available; well-understood convergence theory; mature software implementations.

Cons: converge to the nearest local minimum; performance degrades when the cost landscape contains multiple disconnected basins or non-smooth features. As shown in Chapter 7, the coupled oscillator problem at strong coupling exhibits precisely this structure, and L-BFGS-B fails to find the global optimum from any of the initializations we tested.

3.1.2 Evolutionary and Global Stochastic Methods

Derivative-free evolutionary methods such as CMA-ES [38] maintain a distribution over candidate solutions and adapt its covariance to track the local landscape geometry. They are designed for non-convex and multimodal problems and have been successful on many control benchmarks.

Pros: no gradient required; adaptive to local landscape; in principle capable of escaping local minima.

Cons: require many function evaluations (often 10^4 to 10^6); convergence to the global optimum is not guaranteed; performance depends on hyperparameter choices. In Chapter 7 we demonstrate that CMA-ES with up to 1.96×10^5 function evaluations does not discover the off-cycle optimum.

3.1.3 Pontryagin Shooting and Indirect Methods

Indirect methods solve the necessary conditions from Pontryagin’s minimum principle [39] by reformulating the optimal control problem as a two-point boundary value problem in the state and costate variables.

Pros: provide exact necessary conditions; for low-dimensional problems with smooth Hamiltonians they can yield very accurate solutions.

Cons: the costate dynamics are typically unstable; small initial-guess errors amplify; convergence is sensitive to the shooting initialization. For coupled oscillator problems with multiple cost basins, shooting methods inherit the same locality limitation as gradient methods.

3.2 Reduced-Order Methods

The dominant family of methods for oscillator control exploits the limit-cycle structure to reduce the problem dimension before optimization. The phase variable parameterizes position along the cycle, and additional coordinates capture transverse dynamics with increasing fidelity.

3.2.1 Phase Reduction with iPRC

Standard phase reduction [6, 7] replaces the full state with the phase ϕ and uses the iPRC $Z(\phi)$ to characterize the response to weak inputs. Optimal control problems reduce to one-dimensional variational problems amenable to Pontryagin shooting [8, 9].

Pros: dramatic dimensionality reduction; closed-form solutions in some cases; elegant theoretical framework.

Cons: valid only in the weak-input limit; cannot represent off-cycle trajectories;

saturates well short of strong-coupling targets, as demonstrated in Chapter 7.

3.2.2 Phase Reduction with Interaction Function

For weakly coupled oscillators, the dynamics admit a further reduction to a phase-difference equation with interaction function $H(\psi)$ [6, 25]. Optimal control formulations on this two-dimensional reduction inherit the structure of standard phase reduction.

Pros: captures coupling effects within the reduced-order framework; computationally inexpensive.

Cons: valid only in the weak-coupling regime; saturates near but not at the antiphase target for coupling strengths used in practice.

3.2.3 Phase-Amplitude and Isostable Methods

Phase-amplitude reductions [10, 30, 12] extend the framework with additional coordinates capturing transverse dynamics. Adaptive variants [32, 13] can handle moderately strong inputs.

Pros: valid beyond the weak-perturbation regime; recover transverse dynamics near the cycle; tractable for moderate-amplitude excursions.

Cons: still assume the trajectory lies in a tubular neighborhood of the cycle; no reviewed reduced-order control study certifies a strategy involving excursions of the depth identified in Chapter 7.

3.3 Full-State Methods

Full-state methods compute the optimal control without reducing the dimension of the state space. Unlike direct optimization methods, which search in the finite-dimensional

control space and return a single trajectory, full-state methods solve for the value function on the entire state space and return a feedback law defined at every point. The two principal approaches are HJB-based dynamic programming and sample-based dynamic programming.

3.3.1 Grid-Based HJB Solvers

Grid-based HJB methods discretize the state space and solve the HJB partial differential equation backward in time. The result is a value function $\mathcal{V}(z, t)$ from which the optimal feedback control is derived pointwise [39, 40].

Pros: compute the globally optimal feedback law; no locality assumption; the value function exposes the global structure of the control landscape; the gradient of $\mathcal{V}$ is the policy, so no additional optimization step is required once $\mathcal{V}$ is computed.

Cons: subject to the curse of dimensionality; in practice limited to four to six state dimensions with current grid resolutions and hardware. Tensor decomposition methods [41] and neural HJB methods [42] have been proposed to extend the reach but introduce their own approximation tradeoffs.

3.3.2 Sample-Based and Path-Integral Methods

Path-integral and sample-based dynamic programming methods approximate the value function via Monte Carlo sampling. They scale better with dimension but provide weaker convergence guarantees and introduce sampling variance.

3.3.3 Neural HJB and Deep BSDE

Recent work has approached high-dimensional HJB through neural network parameterization of the value function (PINNs) [42] and through forward-backward SDE formulations

(Deep BSDE) [43]. These methods have demonstrated scaling to dimensions where grid-based methods are infeasible, at the cost of weaker theoretical guarantees.

3.4 Method Selection: The Problem This Thesis Addresses

The methods surveyed above span a spectrum of computational cost and structural assumption. A natural question is which method is appropriate for a given problem instance. The existing literature does not answer this question constructively: reduced-order validity is typically characterized asymptotically (in the limit of weak coupling or weak perturbation) or empirically (by comparison with a reference solution on test cases). Chapter 7 establishes that on the coupled-oscillator problem treated here, the empirical comparison places gradient methods, global stochastic search, and phase-reduced control on the same side of a common failure boundary.

Chapter 4

Why Solve the Hamilton-Jacobi-Bellman Equation in Four Dimensions

This chapter motivates the methodological choice underlying the thesis: solving the full four-dimensional Hamilton-Jacobi-Bellman equation for two coupled oscillators rather than working with a reduced-order or direct-optimization formulation. The chapter also develops the HJB equation, its Hamiltonian decomposition, and the closed-form saturated feedback law in detail. The numerical methods used to solve the resulting PDE on the GPU are deferred to Chapter 5.

4.1 Statement of the Problem

We consider two diffusively coupled Stuart-Landau-like oscillators whose dynamics are

$$\dot{x}_i = ax_i - \omega y_i - ax_i^3 - (a + 2b)x_i y_i^2 + c(x_j - x_i) + u_i, \quad (4.1)$$

$$\dot{y}_i = \omega x_i + ay_i - ay_i^3 + (2b - a)x_i^2 y_i, \quad (4.2)$$

for $i \in \{1, 2\}$, $j \neq i$, with $a > 0$, $\omega > 0$, $b \in [0, \omega)$, $c \geq 0$, and $u_i \in [-u_{\max}, u_{\max}]$.

Control and coupling enter only the x -channel, modeling electrode-based actuation and gap-junction coupling in neural applications [22]. The state $z = (x_1, y_1, x_2, y_2) \in \mathbb{R}^4$.

The cost functional is

$$J(z_0) = \int_0^{T_f} R(u_1^2 + u_2^2) dt + \frac{\gamma}{2} d^2(z(T_f)), \quad (4.3)$$

with terminal cost $d^2(z) = (x_1 + x_2)^2 + (y_1 + y_2)^2 + (r_1 - 1)^2 + (r_2 - 1)^2$ targeting the antiphase orbit, where $r_i^2 = x_i^2 + y_i^2$.

4.1.1 Choice of Parameter Values

The canonical parameter values fixed throughout the thesis are $a = \omega = 1$, $b = 0.8$, $u_{\max} = 1$, $R = 10$, $\gamma = 1000$, and $T_f = 2T$ where T is the natural period of the unit-circle limit cycle. These choices warrant explicit justification:

$a = \omega = 1$. Setting the Floquet exponent magnitude $|\lambda| = 2a$ and the bare frequency ω both to unity nondimensionalizes time so that the unit-circle limit cycle is the natural reference and all other quantities scale relative to it. With these choices, one time unit corresponds to $1/\omega$ in the original units and the Floquet exponent becomes $\lambda = -2$ per time unit.




Figure 4.1: The autonomous Stuart-Landau system at the canonical parameters $a = \omega = 1$, $b = 0.8$. Left: trajectories starting both inside and outside the unit circle spiral onto the limit cycle $r = 1$. Right: angular velocity $\dot{\phi}$ on the cycle for several values of the nonisochronicity parameter b . At $b = 0$ the velocity is constant at $\omega = 1$. For $b > 0$ the velocity oscillates, and its time-average $\Omega = \sqrt{\omega^2 - b^2}$ is strictly less than ω . The difference $\delta\omega = \omega - \Omega$ is the phase-velocity surplus that underlies the amplitude-collapse mechanism in Chapter 7.

$b = 0.8$. Chosen from the sweep over nonisochronicity in Table 7.1 (Chapter 7). At $b = 0.8$, the system exhibits strong nonisochronicity ($\delta\omega = b^2/(\omega + \Omega) = 0.4$) but the natural period remains short enough ($T \approx 10.47$ time units) that on-cycle steering cannot compete with amplitude collapse at moderate coupling. At $b = 0.9$ the period stretches to $T = 14.4$ time units, giving on-cycle steering enough time within the horizon and delaying the collapse transition; at $b = 0.5$ the energetic advantage of collapse is smaller. $b = 0.8$ is the cleanest demonstration regime.

$u_{\max} = 1$. Matches the radius of the unit-circle limit cycle. With this choice, control authority is comparable to the natural oscillator amplitude, which is the physically meaningful scale.

$R = 10$, $\gamma = 1000$. Selected from a 5×3 sweep over $R \in \{1, 10, 25, 50, 100\}$ and $\gamma \in \{500, 1000, 2000\}$. The chosen values balance control authority against value-function regularity: smaller R produces a value function with sharper kinks where the optimal control saturates (harder to resolve numerically), larger R produces a smoother value

function but at the cost of weaker control authority and worse target tracking.

$T_f = 2T$. Two natural periods, approximately 20.94 time units. This horizon gives the controller enough time to execute the three-phase amplitude-collapse strategy: collapse (roughly 5 time units), hold (up to 6 time units at $b = 0.8$), and rebuild (approximately 1.1 time units, set by the Floquet rate). Shorter horizons make the collapse strategy infeasible because the controller cannot harvest enough free phase to justify the hold cost; longer horizons inflate the running cost without changing the optimal mechanism.

4.2 Why Not Phase Reduction Alone

Phase reduction is the natural first attempt. The reduction replaces the four-dimensional state with two phase variables (ϕ_1, ϕ_2) on the limit cycles, yielding the two-dimensional dynamics

$$\dot{\phi}_i = \omega + br^2 \sin 2\phi_i + Z_x(\phi_i) [c(x_\gamma(\phi_j) - x_\gamma(\phi_i)) + u_i], \quad (4.4)$$

where $Z_x(\phi)$ is the iPRC and $x_\gamma(\phi)$ is the on-cycle x -component.

This reduction is appropriate when the coupling and control are weak enough that $z(t)$ remains close to the limit cycle for all t . However, the reduction is *structural*: the state space of the reduced model is the two-torus $\mathbb{T}^2$, not $\mathbb{R}^4$, and trajectories that deviate from the cycle have no representation in the reduced model. If the optimal control mechanism involves such deviation, phase reduction cannot recover it even with arbitrarily fine numerical resolution of the reduced dynamics.

Chapter 7 demonstrates that this is precisely what happens above a critical coupling strength. The minimum amplitude along the optimal trajectory drops from $r_{1,\min} > 0.85$ at weak coupling to $r_{1,\min} \approx 0.1$ at strong coupling, an excursion of $|r - 1| = 0.89$ that




Figure 4.2: Two phase circles $S^1 \times S^1$ form the 2-torus $\mathbb{T}^2$, the natural state space of the phase-reduced two-oscillator system. Phase-reduced optimal control searches over trajectories on this surface. The full-state Hamilton-Jacobi-Bellman formulation searches over $\mathbb{R}^4$ instead, allowing trajectories that leave the limit-cycle neighborhood altogether. Chapter 7 demonstrates that the globally optimal control trajectory at strong coupling does precisely this.

lies well outside the validity range of any phase-based reduction.

4.3 Why Not Direct Optimization in $\mathbb{R}^4$

A natural alternative is to retain the full $\mathbb{R}^4$ state space but parameterize the control $u(t)$ and minimize (4.3) over the parameters using a direct optimization method. This approach avoids the structural restriction of phase reduction while remaining computationally tractable.

The difficulty is that the cost landscape for the coupled oscillator problem at strong coupling contains multiple disconnected basins. The two locally optimal strategies, an on-cycle steering strategy and an off-cycle collapse strategy, are separated by a cost ridge in the space of control parameterizations. Both gradient-based methods (L-BFGS-B) and global stochastic methods (CMA-ES with $> 10^5$ function evaluations) converge to the

on-cycle basin from all initializations we tested. The collapse basin contains the globally optimal solution but is not discoverable by either method in practice.

Chapter 7 reports the detailed comparison. The point for the present chapter is that direct optimization does not, as a matter of empirical fact, find the globally optimal solution for this problem.

4.4 Two Approaches in Comparison

The contrast between direct optimization and dynamic programming is worth stating precisely, since the two paradigms differ in what they solve for and where they search.

Direct optimization searches in control space. The decision variable is a finite-dimensional vector $\mathbf{u} \in \mathbb{R}^N$ representing a discretized control signal (for example piecewise constant on N intervals). The optimization problem is

$$\mathbf{u}^* = \arg \min_{\mathbf{u} \in [-u_{\max}, u_{\max}]^N} J(\mathbf{u}; z_0), \quad (4.5)$$

which is a finite-dimensional minimization. The output is a single trajectory associated with one initial state z_0 . Cost per query is modest (one rollout per cost evaluation) but the global guarantee is absent: gradient methods converge to local minima, and population-based search has the same blind spot when basins are disconnected.

Dynamic programming searches in state space. The decision variable is the value function $\mathcal{V} : \mathbb{R}^4 \times [0, T_f] \rightarrow \mathbb{R}$,

$$\mathcal{V}(z, t) = \min_{u(\cdot)} \int_t^{T_f} R(u_1^2 + u_2^2) ds + \frac{\gamma}{2} d^2(z(T_f)), \quad (4.6)$$

defined at every state and every time. The output is a feedback law $u^*(z, t)$ defined everywhere on the state space, not just along one trajectory. Cost is high (one PDE solve

on G^4 grid points) but the result is globally optimal by Bellman's principle.

The choice between these paradigms is a tradeoff between cost per problem instance and global guarantees. For the coupled oscillator problem, the deterministic results in Chapter 7 show that the global guarantee provided by HJB is essential: direct optimization, in all the configurations we tested, fails to find the globally optimal strategy. This is what motivates the computational investment in solving the 4D HJB equation.

4.5 The Hamilton-Jacobi-Bellman Equation

The value function $\mathcal{V}(z, t)$ defined above satisfies a first-order nonlinear PDE on $[0, T_f] \times \mathbb{R}^4$. Introducing backward time $\tau = T_f - t$, the equation reads

$$\frac{\partial \mathcal{V}}{\partial \tau} + H(z, \nabla_z \mathcal{V}) = 0, \quad \mathcal{V}(z, 0) = \frac{\gamma}{2} d^2(z), \quad (4.7)$$

where the initial condition at $\tau = 0$ corresponds to physical time $t = T_f$, i.e. the terminal cost in the original problem.

4.5.1 Hamiltonian Decomposition

The Hamiltonian admits an explicit decomposition into a drift contribution and a control contribution. Writing $p = \nabla_z \mathcal{V}$ for the value-function gradient, the Hamiltonian is

$$H(z, p) = -f(z) \cdot p + H^*(p_{x_1}) + H^*(p_{x_2}), \quad (4.8)$$

where $f(z)$ is the autonomous drift from (4.1)–(4.2) (i.e. the right-hand side without the control terms), and H^* is the convex dual of the running cost on each control channel.

The first term, $-f(z) \cdot p$, is the drift Hamiltonian. It captures the contribution to

the cost-to-go from the autonomous dynamics. The second and third terms encode the optimization over u_1 and u_2 , which has been carried out analytically in closed form using the convex-dual structure.

4.5.2 The Convex Dual and the Saturated Feedback Law

For running cost Ru_i^2 with bounded control $|u_i| \leq u_{\max}$, the optimization $\min_u \{Ru^2 + u \cdot p\}$ on the constraint set is piecewise quadratic. The minimizer is

$$u_i^* = \text{sat}\left(-\frac{p_{x_i}}{2R}, u_{\max}\right), \quad (4.9)$$

where $\text{sat}(\xi, \xi_{\max}) = \text{sign}(\xi) \min(|\xi|, \xi_{\max})$ denotes saturation at $\pm\xi_{\max}$. Substituting back, the convex dual is

$$H^*(p) = \begin{cases} \frac{p^2}{4R}, & |p| \leq 2Ru_{\max}, \\ -Ru_{\max}^2 + u_{\max} |p|, & |p| > 2Ru_{\max}. \end{cases} \quad (4.10)$$

The piecewise structure has a clear interpretation. When $|p|$ is small, the unconstrained optimum lies inside the constraint set and the dual is the standard quadratic form. When $|p|$ is large, the constraint is active, the optimal control saturates at $\pm u_{\max}$, and the dual becomes linear in $|p|$.

4.5.3 What Full-State HJB Provides

Two features of this approach are essential for the thesis.

First, the HJB solution is globally optimal by construction. The off-manifold mechanism that direct optimization fails to find is, in contrast, present in the HJB feedback law automatically. We do not need to know the mechanism in advance to find it; it emerges

from the computation. The gradient of $\mathcal{V}$ is literally the policy, in the sense that the feedback law is recovered pointwise from $-\nabla_x \mathcal{V}$ by the formula above with no additional optimization step.

Second, the value function exposes the global structure of the control landscape, including the basin separation that explains the failure of direct optimization. The HJB approach is thus both a tool for finding the optimum and a diagnostic for why other methods fail.

4.6 Cost: Four Dimensions Is the Boundary of Feasibility

The principal drawback of grid-based HJB is the curse of dimensionality. A grid with G points per dimension requires G^d storage and a comparable number of operations per time step. For $G = 64$ and $d = 4$, this is approximately 1.7×10^7 grid points, which fits comfortably in modern GPU memory. For $d = 6$ (three oscillators) the same grid resolution gives approximately 6.9×10^{10} grid points, which exceeds practical limits.

Four dimensions is therefore at the boundary of what grid-based HJB can do with current hardware. The two-oscillator problem is the largest coupled-oscillator system for which full-state HJB provides a definitive answer, which makes it the right setting to identify mechanisms that may then be sought, by other means, in larger systems.

Chapter 5

Numerical Methods for the Hamilton-Jacobi-Bellman Equation

This chapter describes the numerical methods used to solve the HJB equation (4.7) on the four-dimensional state space. The discretization is structured as a method-of-lines approach: spatial discretization of the Hamiltonian via a high-order WENO scheme produces a system of ordinary differential equations in time, which is integrated with a strong-stability-preserving Runge-Kutta scheme. Section 5.4 gives the explicit solver pseudocode for one backward step.

5.1 The Hamiltonian and Optimal Feedback

The Hamiltonian decomposition (4.8) and the convex dual (4.10) from Chapter 4 are repeated here for convenience:

$$H(z, p) = -f(z) \cdot p + H^*(p_{x_1}) + H^*(p_{x_2}), \quad (5.1)$$

with

$$H^*(p) = \begin{cases} \frac{p^2}{4R}, & |p| \leq 2Ru_{\max}, \\ -Ru_{\max}^2 + u_{\max}|p|, & \text{otherwise,} \end{cases} \quad (5.2)$$

yielding the saturated feedback law $u_i^* = \text{sat}(-p_{x_i}/2R, u_{\max})$.

The piecewise structure of H^* is the source of the kinks in $\mathcal{V}$ that motivate the WENO discretization below. Wherever $|p_{x_i}| = 2Ru_{\max}$ in state space, H^* has a discontinuous second derivative, and the corresponding viscosity solution $\mathcal{V}$ has a kink in its gradient. A high-order linear scheme would oscillate violently near these kinks; WENO5 handles them gracefully by adaptive stencil weighting.

5.2 Spatial Discretization:

Local Lax-Friedrichs WENO5

The Hamiltonian $H(z, \nabla \mathcal{V})$ is discretized using a fifth-order weighted essentially non-oscillatory (WENO5) scheme [44] with local Lax-Friedrichs flux splitting. WENO schemes provide high-order accuracy in smooth regions while avoiding spurious oscillations near non-smooth features, which is essential for HJB because viscosity solutions can have kinks where the optimal feedback saturates or switches.

5.2.1 WENO5 Reconstruction

For a one-dimensional derivative $\partial \mathcal{V} / \partial z$ approximated at grid point i , WENO5 uses three candidate third-order stencils that share grid points near i and combines them with

adaptive nonlinear weights ω_k , $k \in \{0, 1, 2\}$, yielding

$$\left(\frac{\partial \mathcal{V}}{\partial z}\right)_i^\pm = \omega_0^\pm q_0^\pm + \omega_1^\pm q_1^\pm + \omega_2^\pm q_2^\pm, \quad (5.3)$$

where $q_k^\pm$ are the three candidate reconstructions from biased stencils and the weights $\omega_k^\pm$ depend on smoothness indicators computed from the local solution. In smooth regions all three stencils are equally informative and the weights converge to the linear-optimal values that give fifth-order accuracy; near a kink, one or more candidates exhibit large smoothness indicators and the weights shift toward the smoothest candidate, reducing to third-order accuracy locally but preserving non-oscillatory behavior.

5.2.2 Local Lax-Friedrichs Flux

For a Hamiltonian $H(p)$ depending on the gradient p approximated by left and right WENO reconstructions $p_i^\pm$, the Lax-Friedrichs numerical Hamiltonian is

$$\hat{H}(p^+, p^-) = H\left(\frac{p^+ + p^-}{2}\right) - \frac{\alpha}{2}(p^+ - p^-), \quad (5.4)$$

where $\alpha = \max_p |H'(p)|$ is the local viscosity coefficient computed from the maximum wave speed in each direction. The scheme is monotone, consistent, and Lipschitz, which are the three conditions required by the Barles-Souganidis theorem [45] for convergence to the viscosity solution. In the four-dimensional setting this construction is applied dimension by dimension.

5.3 Time Integration: SSP-RK3

Backward-in-time integration is performed with the third-order strong-stability-preserving Runge-Kutta scheme [46]. Writing $\mathcal{L}(\mathcal{V})$ for the spatial operator (i.e. $\mathcal{L}(\mathcal{V}) = -H(z, \nabla^{\text{WENO5}}\mathcal{V})$ with the WENO5 gradient reconstruction and the LLF numerical Hamiltonian), the three-stage SSP-RK3 update is

$$\mathcal{V}^{(1)} = \mathcal{V}^n + \Delta\tau \mathcal{L}(\mathcal{V}^n), \quad (5.5)$$

$$\mathcal{V}^{(2)} = \frac{3}{4}\mathcal{V}^n + \frac{1}{4}\mathcal{V}^{(1)} + \frac{1}{4}\Delta\tau \mathcal{L}(\mathcal{V}^{(1)}), \quad (5.6)$$

$$\mathcal{V}^{n+1} = \frac{1}{3}\mathcal{V}^n + \frac{2}{3}\mathcal{V}^{(2)} + \frac{2}{3}\Delta\tau \mathcal{L}(\mathcal{V}^{(2)}). \quad (5.7)$$

Each stage is a convex combination of forward-Euler steps, which preserves the non-oscillatory property of WENO5 across the time integration. The CFL constraint is

$$\Delta\tau \leq \text{CFL} \cdot \frac{\Delta x}{\max_z \alpha(z)}, \quad (5.8)$$

where Δx is the grid spacing and α is the local Lax-Friedrichs viscosity. For the Stuart-Landau problem with the canonical parameters, the resulting time step gives approximately 3,000 backward steps to integrate from $\tau = 0$ to $\tau = T_f$.

5.4 Solver Pseudocode

The full backward integration is summarized as Algorithm 1. The parallel structure is over grid points: each iteration of the inner loop is independent, which is what makes the GPU implementation efficient (Chapter 6).

Algorithm 1 One SSP-RK3 backward step on the GPU. Each stage update is applied pointwise, with one CUDA thread per grid point.

```

1: Initialization:  $\mathcal{V}^0(z) \leftarrow \frac{\gamma}{2} d^2(z)$  on grid  $G^4$  with  $G = 64$ 
2: for each backward step  $\tau_n \rightarrow \tau_{n+1}$  with CFL-limited  $\Delta\tau$  do
3:    $\mathcal{V}^{(1)} \leftarrow \mathcal{V}^n + \Delta\tau \mathcal{L}(\mathcal{V}^n)$  ▷ stage 1
4:    $\mathcal{V}^{(2)} \leftarrow \frac{3}{4}\mathcal{V}^n + \frac{1}{4}\mathcal{V}^{(1)} + \frac{1}{4}\Delta\tau \mathcal{L}(\mathcal{V}^{(1)})$  ▷ stage 2
5:    $\mathcal{V}^{n+1} \leftarrow \frac{1}{3}\mathcal{V}^n + \frac{2}{3}\mathcal{V}^{(2)} + \frac{2}{3}\Delta\tau \mathcal{L}(\mathcal{V}^{(2)})$  ▷ stage 3
6: end for
7: return  $\mathcal{V}^{N_\tau}$ 

```

Spatial operator $\mathcal{L}(\mathcal{V})$ at each grid point:

```

8: Reconstruct  $(\partial_{x_1}\mathcal{V})^\pm, (\partial_{y_1}\mathcal{V})^\pm, (\partial_{x_2}\mathcal{V})^\pm, (\partial_{y_2}\mathcal{V})^\pm$  via WENO5
9: Form  $\hat{H}_{\text{LLF}}$  using (5.4) with local wave speeds  $\alpha_{x_i}, \alpha_{y_i}$ 
10: return  $-\hat{H}_{\text{LLF}}(z, \nabla^{\text{WENO5}}\mathcal{V})$ 

```

5.5 Boundary Conditions

The computational domain $[-1.5, 1.5]^4$ is chosen to contain all trajectories of interest with margin. Boundary conditions are set by linear extrapolation of the value function, which is appropriate for outflow boundaries and does not introduce spurious information into the interior. The implementation also supports constant, quadratic extrapolation, and periodic boundary conditions for the other validated problems (Chapter 6).

5.6 Verification: What Each Test Proves

The numerical solution is verified through three convergence checks, each probing a different aspect of the solver. It is worth stating explicitly what each test does and does not certify, because the three tests are not equivalent and naively presenting them as “three convergence checks” obscures the distinct role of each.

5.6.1 Test 1: Linear Advection in 4D (the Fifth-Order Test)

Linear advection on a four-dimensional periodic domain admits a smooth exact solution. We measure the empirical convergence order as the grid is refined from $G = 16$ to $G = 64$ and obtain measured orders between 4.84 and 4.99 in both L^2 and L^∞ . This is the test that genuinely probes the fifth-order accuracy of WENO5: because the exact solution is smooth, the WENO5 weights stay near the linear-optimal values, and the fifth-order accuracy shows up in the error decay rate. A lower-order scheme would give a measured order strictly below four; a fifth-order scheme gives a measured order at most five.

5.6.2 Test 2: 4D LQR vs Analytical Riccati (End-to-End Pipeline Check)

For a four-dimensional double integrator with quadratic terminal and running costs, the value function satisfies an algebraic Riccati equation whose analytical solution is computable. Solving the corresponding HJB equation numerically and comparing to the analytical value function gives a relative L^2 error of 6.3×10^{-5} at $G = 64$.

This is a powerful test of the end-to-end pipeline, including the Hamiltonian implementation, the boundary handling, the time integration, and the optimal control extraction. However, it is *not* a fifth-order convergence test. The LQR value function is exactly quadratic in the state, $\mathcal{V}(z) = z^\top P z$ for the Riccati matrix P , and a quadratic is captured exactly by any sufficiently high-order scheme on a fine enough grid. A third-order scheme, a fourth-order scheme, and a fifth-order scheme would all converge to the analytical Riccati solution at machine precision on the LQR problem; the test cannot distinguish them. What the LQR test *does* prove is that the solver as a whole produces the correct optimal control on a problem where the answer is known analytically.

5.6.3 Test 3: Stuart-Landau Grid Convergence

(Self-Convergence on the Real Problem)

For the actual problem of interest, where no analytical answer exists, we refine from $G = 64$ to $G = 80$ to $G = 96$ and measure the change in the optimal cost and the minimum amplitude along the optimal trajectory. The cost variation is below 0.5% across the three resolutions, and the optimal mechanism (amplitude collapse or on-cycle steering) is invariant. This test confirms convergence on the actual problem, complementing the order check on linear advection and the analytical check on LQR.

5.6.4 Domain Extension and Symmetry

In addition to the three convergence tests, two further checks confirm robustness. First, enlarging the computational domain from $[-1.5, 1.5]^4$ to $[-2.0, 2.0]^4$ changes the optimal cost by less than 0.3%, confirming that the chosen domain is large enough not to constrain the solution. Second, the system is $\mathbb{Z}_2$ -equivariant under oscillator swap [47], and the computed value function satisfies this symmetry to within discretization error.

5.6.5 Summary

Putting these together: WENO5's fifth-order accuracy is established by Test 1 (linear advection); the full pipeline's correctness against a known analytical solution is established by Test 2 (LQR vs Riccati); and convergence on the actual Stuart-Landau problem is established by Test 3 (grid refinement). The Barles-Souganidis theorem [45] ties these together: a monotone, consistent, stable scheme converges to the viscosity solution, and our scheme satisfies all three conditions by construction.

Chapter 6

GPU Acceleration and the Extension of CASL-HJX to Four Dimensions

This chapter describes the GPU-accelerated implementation that makes four-dimensional HJB solutions practical, and the extension of the CASL-HJX framework [14] from two to four spatial dimensions.

6.1 Computational Cost of 4D HJB

A grid with $G = 64$ points per dimension on a four-dimensional state space contains $G^4 \approx 1.68 \times 10^7$ grid points. Each backward-time step requires evaluation of the WENO-reconstructed Hamiltonian at every grid point in every dimension, followed by an SSP-RK3 update. Storage requirements include the value function, intermediate RK stages, and reconstruction buffers, totaling approximately 0.5 GB in single precision and 1.0 GB in double precision per copy. With several working copies for the RK stages this fits comfortably within the 40 GB memory of an NVIDIA A100.

A representative HJB solve on $[-1.5, 1.5]^4$ over $T_f = 2T$ with the canonical parameters

$a = \omega = 1$, $b = 0.8$, $u_{\max} = 1$ completes in approximately five minutes on a single A100. A comparable CPU implementation would require an estimated 10 to 20 hours on a modern multicore workstation.

6.2 Parallelization Strategy

The HJB update is structurally parallel in space: at each time step, the WENO reconstruction and Hamiltonian evaluation at one grid point depend only on a local stencil of neighbors. The parallel structure maps naturally to a CUDA grid in which each thread handles one or a small block of spatial grid points.

Memory layout. The value function is stored in a contiguous four-dimensional array with the innermost axis chosen to maximize coalesced memory access for the WENO stencil. Stencils along the other three axes incur strided access; this is handled either by axis-by-axis transposed sweeps or by sufficient on-chip caching depending on the relative cost in each direction.

Kernel structure. A separate CUDA kernel handles the WENO reconstruction along each axis. The kernel reads the stencil values from global memory, performs the WENO smoothness-indicator computation and the nonlinear weighting, and writes the reconstructed left- and right-biased derivative approximations to a working buffer. A second kernel combines these to produce the local Lax-Friedrichs numerical Hamiltonian and computes the time-derivative contribution.

SSP-RK3 stages. The three RK stages are implemented as successive kernel calls, each performing a weighted combination of previous stages. Synchronization between stages is handled implicitly by the CUDA stream serialization.

6.3 Extension from 2D to 4D

The CASL-HJX framework [14] provides a CPU and GPU implementation of HJB solvers for two-dimensional problems. The extension to four dimensions required three principal changes.

Stencil generalization. The WENO reconstruction was generalized from two-dimensional to four-dimensional grids. The reconstruction itself is one-dimensional along each axis and was kept algorithmically unchanged, but the indexing infrastructure and kernel launch configurations were rewritten to support arbitrary axis selection in four dimensions.

Memory management. The two-dimensional implementation could keep multiple working copies of the value function in shared memory or cache. The four-dimensional case requires explicit management of global-memory buffers because the working set exceeds on-chip storage. The implementation uses pinned host memory for fast host-to-device transfer of initial and boundary data.

6.4 Performance Profile

Profiling on an A100 shows the following approximate breakdown for a deterministic 4D HJB solve at $G = 64$:

- WENO reconstruction: $\sim 55\%$ of runtime, dominated by global memory traffic.
- Numerical Hamiltonian and RK update: $\sim 20\%$, compute-bound.
- Boundary handling and bookkeeping: $\sim 10\%$.
- Time-step adaptation and CFL computation: $\sim 5\%$.
- Output and checkpointing: $\sim 10\%$.

6.5 Code Availability and Software Architecture

The four-dimensional GPU solver is implemented in approximately 7,100 lines of C++/CUDA, structured as a continuation of the two-dimensional CASL-HJX codebase [14]. The repository is organized to separate shared numerical infrastructure from problem-specific Hamiltonian kernels, so that a new four-dimensional optimal control problem can be added by replacing only the Hamiltonian kernel without modifying the solver core.

6.5.1 Repository Structure

The codebase is divided into three layers:

Shared infrastructure (`common/`). Templated four-dimensional array classes with explicit index calculation and boundary padding (`CaslArray4D`), grid descriptors (`CaslGrid4D`), the base Hamiltonian interface (`CaslHamiltonian4D`), an RK4 forward integrator for trajectory reconstruction (`CaslRK4Integrator4D`), and post-processing utilities including quadrilinear interpolation of saved gradient fields. These components handle anything that does not depend on the specific dynamics.

GPU solver kernel (`gpu_opt/casl_gpu_solver4D.cuh`). A single 423-line header containing the WENO5 reconstruction kernels, the local Lax-Friedrichs flux computation, the SSP-RK3 time integration, and four boundary-condition implementations (constant, linear, quadratic extrapolation, and periodic). The WENO5 kernels are direct CUDA translations of the verified CPU implementations from CASL-HJX, with identical index formulas and identical numerical behavior. Boundary conditions are selected through an integer parameter at solver initialization, allowing different conditions on different axes without recompilation.

Problem-specific Hamiltonian kernels (`projects/`). Five problems are imple-

mented, each in a separate subdirectory: coupled Stuart-Landau oscillators (`StuartLandau4D/`, the primary application), coupled Van der Pol oscillators (`VdPDuffing4D/`, the falsification check from Chapter 7), a four-dimensional double integrator with analytical algebraic Riccati equation solution (`LQR4D/`, for verification), four-dimensional linear advection with periodic boundary conditions (`LinearAdvection4D/`, for WENO5 convergence order verification), and a coupled Hodgkin-Huxley reduction (`CoupledNeurons4D/`, exploratory work on neural population control). Each problem provides a Hamiltonian device function and a maximum-wave-speed function in approximately three hundred lines of CUDA, with no modifications required to the solver core.

6.5.2 Templated Solver Pattern

The GPU solver is templated on a problem-specific Hamiltonian struct. For the Stuart-Landau problem the struct exposes a `H()` device function that evaluates the optimal Hamiltonian at each grid point, together with four `maxAbsH*()` device functions that return the maximum wave speed along each axis (needed for the local Lax-Friedrichs flux splitting). The solver is otherwise unmodified across problems, which means that the WENO5 reconstruction, the SSP-RK3 update, the boundary handling, and the time-step adaptation logic are verified in one place and reused everywhere.

This pattern allowed the addition of the Van der Pol falsification system and the LQR verification system, each requiring approximately one day of work, by writing only the Hamiltonian and the dynamics evaluator.

6.5.3 Numerical Validation Suites

The repository includes three numerical validation suites that each test a separate component of the solver, corresponding to the three convergence checks discussed in

Section 5.6. The linear advection problem confirms 5th-order convergence of the WENO5 reconstruction on a smooth solution with periodic boundary conditions (measured orders 4.84 to 4.99 in both L^2 and L^∞). The LQR problem confirms agreement with the analytical solution of the algebraic Riccati equation on a four-dimensional double integrator (relative L^2 error of 6.3×10^{-5} at $G = 64$). The Stuart-Landau problem confirms grid convergence of the value function and the optimal cost across $G \in \{32, 48, 64, 80, 96\}$.

6.5.4 Hardware and Performance

The solver targets NVIDIA GPUs with compute capability 7.5 or higher (T4, A100, H100). Single-precision execution fits the $G = 64$ Stuart-Landau problem in 1.3 GB of VRAM in approximately five minutes; double-precision execution doubles the VRAM and adds roughly 40% to the wall time. The $G = 96$ problem uses 6.8 GB and runs in approximately twenty-five minutes. Beyond $G = 96$ the four-dimensional memory footprint exceeds the 40 GB capacity of an A100, which sets the practical resolution ceiling for the current implementation. Adaptive mesh refinement or tensor-train compression would be required to reach higher effective resolution; this is identified as a future direction in Chapter 8.

The development workflow used rented A100 instances on a commercial GPU cloud (Vast.ai) supplemented by the UCSB Pod cluster. The full 36-point (c, b) parameter sweep reported in Chapter 7 required approximately three hours of A100 time. Repository access and configuration details are documented in the included README; the repository is currently private and will be released publicly with the associated journal publication.

Chapter 7

Deterministic Results: Amplitude Collapse and Multi-Method Failure

This chapter presents the deterministic results that form the core contribution of the thesis. The content corresponds to the work submitted to CDC 2026 [15], adapted and expanded for the thesis. We identify a sharp strategy transition at a critical coupling c^* , characterize the resulting amplitude-collapse mechanism, and demonstrate that this mechanism is inaccessible to gradient-based optimization, global stochastic search, and phase-reduced control.

7.1 Problem Formulation

7.1.1 Two Coupled Stuart-Landau-Like Oscillators

We consider two diffusively coupled oscillators whose dynamics take the Stuart-Landau-like form:

$$\dot{x}_i = ax_i - \omega y_i - ax_i^3 - (a + 2b)x_i y_i^2 + c(x_j - x_i) + u_i, \quad (7.1)$$

$$\dot{y}_i = \omega x_i + ay_i - ay_i^3 + (2b - a)x_i^2 y_i, \quad (7.2)$$

for $i \in \{1, 2\}$, $j \neq i$, with $a > 0$, $\omega > 0$, $b \in [0, \omega)$, $c \geq 0$, and $u_i \in [-u_{\max}, u_{\max}]$. Control and coupling enter only the x -channel, modeling electrode-based actuation and gap-junction coupling in neural applications [22].

We refer to these as “Stuart-Landau-like” because of the extra $\sin 2\phi$ term that appears in the polar dynamics. In polar coordinates $(x_i, y_i) = r_i(\cos \phi_i, \sin \phi_i)$, the uncoupled dynamics are

$$\dot{r} = ar(1 - r^2), \quad \dot{\phi} = \omega + br^2 \sin 2\phi. \quad (7.3)$$

This particular form was chosen because, in addition to the in-phase state being stable, the $\sin 2\phi$ nonisochronicity makes the anti-phase solution with $x_2(t) = -x_1(t)$ and $y_2(t) = -y_1(t)$ stable for the canonical parameter values, providing a natural target state for the control problem. Note, however, that the basin of attraction for the anti-phase solution is small, so a control strategy targeting this solution must be precise. The unit-circle limit cycle has period $T = 2\pi/\Omega$ where $\Omega = \sqrt{\omega^2 - b^2}$, and Floquet exponent $\lambda = -2a$ governing transverse perturbations.

7.1.2 Time-Unit Conventions

With the canonical parameters $a = \omega = 1$, $b = 0.8$ (justified in Section 4.1.1), the relevant time-scale quantities are:

$$\Omega = \sqrt{\omega^2 - b^2} = \sqrt{1 - 0.64} = 0.6,$$

$$T = \frac{2\pi}{\Omega} \approx 10.47 \text{ time units (natural period),}$$

$$T_f = 2T \approx 20.94 \text{ time units (control horizon),}$$

$$\lambda = -2a = -2 \text{ per time unit (Floquet exponent).}$$

The Floquet rate $\lambda = -2$ predicts that transverse perturbations decay as $\exp(-2t)$ per time unit, providing the time scale on which the amplitude rebuilds after a collapse phase. Specifically, a rebuild from $r_{1,\min} = 0.11$ to $r_1 = 0.99$ takes approximately $\ln(89)/2 \approx 1.1$ time units, which matches the empirical rebuild time observed in the figures below.

[Symmetry] The system is $\mathbb{Z}_2$ -equivariant under oscillator swap [47]. The HJB equation inherits this symmetry; we verify it *a posteriori* to within discretization error. The initial condition $z_0 = (1, 0, \cos 0.1, \sin 0.1)$ breaks this symmetry; the controller selects which oscillator to collapse based on the state.

7.1.3 Optimal Control Objective

The cost functional is

$$J(z_0) = \int_0^{T_f} R(u_1^2 + u_2^2) dt + \frac{\gamma}{2} d^2(z(T_f)), \quad (7.4)$$




Figure 7.1: Phase portrait of the autonomous Stuart-Landau system. Trajectories spiral toward the unit-circle limit cycle. Phase velocity along the limit cycle is $\Omega = \sqrt{\omega^2 - b^2}$; phase velocity at the origin is the unperturbed rate ω . The difference $\delta\omega = \omega - \Omega$ provides the phase budget exploited by the off-manifold control mechanism identified in Section 7.5.

with $R = 10$, $\gamma = 1000$, $T_f = 2T$, and terminal cost

$$d^2(z) = (x_1 + x_2)^2 + (y_1 + y_2)^2 + (r_1 - 1)^2 + (r_2 - 1)^2 \quad (7.5)$$

targeting the entire anti-phase orbit. Since d^2 depends on $(x_1 + x_2, y_1 + y_2)$ and radii r_i but not the absolute phase, it vanishes for any anti-phase pair $(re^{i\phi}, re^{i(\phi+\pi)})$ with $r = 1$, regardless of ϕ . The $(r_i - 1)^2$ terms prevent the trivial solution $z = 0$ from achieving zero terminal cost.

7.2 HJB Solution Method

Let $\mathcal{V}(z, \tau)$ denote the value function in backward time $\tau = T_f - t$. It satisfies the Hamilton-Jacobi-Bellman equation

$$\frac{\partial \mathcal{V}}{\partial \tau} + H(z, \nabla_z \mathcal{V}) = 0, \quad \mathcal{V}(z, 0) = \frac{\gamma}{2} d^2(z), \quad (7.6)$$

where the initial condition at $\tau = 0$ corresponds to the terminal cost of the original problem. The Hamiltonian is

$$H(z, p) = -f(z) \cdot p + H^*(p_{x_1}) + H^*(p_{x_2}), \quad (7.7)$$

with $p = \nabla_z \mathcal{V}$ and $f(z)$ the autonomous drift from (7.1)–(7.2). The convex dual H^* , corresponding to the running cost Ru^2 under $|u| \leq u_{\max}$, is

$$H^*(p) = \begin{cases} \frac{p^2}{4R}, & |p| \leq 2Ru_{\max}, \\ -Ru_{\max}^2 + u_{\max}|p|, & \text{otherwise,} \end{cases} \quad (7.8)$$

yielding the feedback control

$$u_i^* = \text{sat}\left(-\frac{p_{x_i}}{2R}, u_{\max}\right). \quad (7.9)$$

The PDE is solved on $[-1.5, 1.5]^4$ using a uniform grid with $G = 64$ points per dimension, using the WENO5+SSP-RK3 scheme described in Chapter 5 and the GPU-accelerated implementation described in Chapter 6, enabling full-state solutions in approximately five minutes on a single NVIDIA A100. All optimal trajectories remain well within the computational domain ($\|z(t)\|_{\infty} < 1.2$), and enlarging the domain to $[-2.0, 2.0]^4$ changes the optimal cost by less than 0.3%. Grid refinement up to $G = 96$ yields cost variations below 0.5%.

7.3 Two Qualitatively Distinct Strategies

The HJB solution reveals two qualitatively distinct optimal strategies depending on the coupling strength c . At weak coupling ($c = 0.10$), both oscillators remain near the

limit cycle throughout the control horizon. The controller applies small perturbations to smoothly steer the phase difference toward 180° . The trajectories in the (x_i, y_i) phase plane trace paths close to the unit circle, and the minimum amplitude stays above $r_{1,\min} > 0.85$. This is consistent with standard phase reduction intuition: the oscillators stay near their limit cycle while the controller adjusts the timing.

The picture changes dramatically at stronger coupling. At $c = 0.30$, the controller drives oscillator 1 to $r_{1,\min} = 0.11$, nearly to the origin, while oscillator 2 remains near the limit cycle. This is clearly visible in the phase portrait: the trajectory of oscillator 1 spirals inward to a small neighborhood of the origin, dwells there while accumulating a phase advantage, then spirals back out to rejoin the limit cycle in approximate anti-phase with oscillator 2. We call this *amplitude collapse*.

The contrast between $c = 0.20$ and $c = 0.30$ is particularly striking. At $c = 0.20$, the trajectory stays near the limit cycle ($r_{1,\min} = 0.77$); at $c = 0.30$, it collapses to $r_{1,\min} = 0.11$. There is no intermediate regime of moderate amplitude excursion: the controller either stays on-cycle or commits fully to collapse. This sharp difference between adjacent coupling values motivates a dense sweep to precisely locate the transition.

7.3.1 Saturating-Optimal Structure of the Control

The control signal at $c = 0.30$ exhibits three visually distinct phases: a large pulse during the collapse, a quiescent interval during the hold, and another large pulse during the rebuild. The qualitative appearance is that of a bang-bang control. The precise characterization, however, is one of saturation rather than strict bang-bang in the Pontryagin sense.

Strict bang-bang control arises when the Hamiltonian is linear in the control on the constraint set, in which case the optimizer always selects one of the bounds $\pm u_{\max}$ with

sharp switches between them. In our problem the running cost is quadratic in u , so the Hamiltonian is quadratic in u as well, and the unconstrained optimal control is the smooth feedback $u^* = -p_{x_i}/(2R)$. The constraint $|u| \leq u_{\max}$ clips this when $|p_{x_i}| > 2Ru_{\max}$, producing the saturated form $u_i^* = \text{sat}(-p_{x_i}/2R, u_{\max})$. The visual signature is “bang-like” because the value-function gradient is large during the collapse and rebuild phases (saturating the unconstrained optimum), and small during the hold (when the natural dynamics together with the vanished diffusive coupling do the controller’s work for free). In this thesis we use the phrase *saturating-optimal* for this structure, reserving *bang-bang* for the strict Pontryagin case.

7.4 Sharp Strategy Transition

A dense sweep over 30 coupling values in the range $c \in [0.01, 0.50]$ reveals a sharp transition at $c^* = 0.227 \pm 0.003$. At $c = 0.2247$, the optimal trajectory stays near the limit cycle with $r_{1,\min} = 0.728$. At $c = 0.2299$, a change of only $\Delta c = 0.005$, the optimal trajectory collapses to $r_{1,\min} = 0.097$. The transition width $\Delta c < 0.006$ indicates a steep crossover between two distinct strategy basins rather than a smooth interpolation. The controller does not gradually increase its off-cycle excursion as coupling increases; instead, it abruptly switches from one strategy to the other.

Strikingly, the optimal cost $J^*(c)$ is *non-monotonic*: it decreases between $c = 0.20$ and $c = 0.30$ despite the stronger coupling opposition. This is a signature of the collapse strategy’s efficiency, which we explain in the next section.

The coupling parameter c is dimensionless under the nondimensionalization $a = \omega = 1$ adopted throughout: c enters the dynamics as a rate $c(x_j - x_i)$ with units of inverse time, but with our time unit set equal to $1/\omega$ the parameter becomes dimensionless. The labels “weak” and “strong” refer to whether c lies below or above the empirical critical coupling

c^* .

7.5 Physical Mechanism: Phase Velocity Surplus

The non-monotonicity of $J^*(c)$ and the efficiency of the collapse strategy are explained by the *phase velocity surplus*. In polar form, $\dot{\phi} = \omega + br^2 \sin 2\phi$. At $r \approx 0$, both the coupling force ($\propto cr_1$) and the nonisochronous distortion ($\propto br^2$) vanish, and the collapsed oscillator rotates uniformly at $\dot{\phi} = \omega$. On the limit cycle ($r = 1$), the time-averaged angular velocity is $\Omega = \sqrt{\omega^2 - b^2}$, which is *slower* than ω whenever $b \neq 0$. The difference

$$\delta\omega = \omega - \Omega = \frac{b^2}{\omega + \Omega} \quad (7.10)$$

is the phase velocity surplus: the collapsed oscillator gains phase at rate $\delta\omega$ relative to oscillator 2, which continues on the limit cycle at rate Ω . At $b = 0.8$, $\omega = 1$: $\delta\omega = 0.4$, providing substantial free phase repositioning without any control effort.

During the hold interval τ (the time oscillator 1 spends at small amplitude), the controller must cancel the coupling force $c \cos \phi_2$ that pushes oscillator 1 away from the origin, at cost $E_{\text{hold}} \approx Rc^2\tau/2$ (since $\langle \cos^2 \phi_2 \rangle = 1/2$). But it gains $\delta\omega \cdot \tau$ radians of free phase. The trade-off between this hold cost ($O(c^2\tau)$) and the phase gained determines whether collapse is energetically favorable. On-cycle control pays a coupling penalty over the entire horizon T_f , while collapse pays only during $\tau \ll T_f$.

The rebuild from $r_{1,\min} = 0.11$ to $r_1 = 0.95$ takes approximately $(2a)^{-1} \ln(0.95/0.11) = 1.1$ time units, less than 11% of T_f ; all computed trajectories satisfy $r_1(T_f) > 0.99$.

Table 7.1: Minimum amplitude $r_{1,\min}$ across the (c, b) parameter plane. Bold entries indicate amplitude collapse ($r_{1,\min} < 0.3$). All runs use $G = 64$, $a = \omega = 1$.

b	$c=0.05$	$c=0.10$	$c=0.20$	$c=0.30$	$c=0.40$	$c=0.50$
0.0	0.928	0.885	0.806	0.181	0.218	0.234
0.2	0.908	0.857	0.112	0.191	0.219	0.223
0.4	0.900	0.858	0.113	0.190	0.207	0.196
0.6	0.901	0.862	0.102	0.177	0.165	0.272
0.8	0.897	0.857	0.773	0.110	0.131	0.184
0.9	0.870	0.845	0.801	0.761	0.337	0.374

7.6 Role of Nonisochronicity

Having established the collapse mechanism at $b = 0.8$, we now investigate how it depends on the nonisochronicity parameter b . Table 7.1 shows $r_{1,\min}$ across 36 (c, b) pairs.

Three findings emerge. First, amplitude collapse occurs at all b values for sufficiently strong coupling, including the isochronous case $b = 0$. The critical coupling c^* is non-monotonic in b : it is lowest at $b = 0.2$ to 0.6 ($c^* \in (0.10, 0.20)$), higher at $b = 0$ ($c^* \in (0.20, 0.30)$), and highest at $b = 0.9$ ($c^* > 0.50$).

Second, while collapse occurs even at $b = 0$, the *energetic advantage* of collapse depends critically on nonisochronicity. At $c = 0.30$, the optimal cost drops from $J^* = 19.93$ at $b = 0$ to $J^* = 11.53$ at $b = 0.8$, a 42% reduction. This is because the phase velocity surplus $\delta\omega$ grows with b : at $b = 0$, $\delta\omega = 0$ and collapse provides no free phase; at $b = 0.8$, $\delta\omega = 0.4$ and the collapsed oscillator gains 138° of free phase repositioning (Table 7.2).

Third, at $b = 0.9$ the long period ($T = 14.4$) provides sufficient time for on-cycle steering, delaying the onset of collapse to higher coupling values.

Table 7.2: Phase budget during collapse ($r_1 < 0.5$) at $c = 0.30$. The free phase $\delta\omega \cdot \tau$ is the phase gained by operating at collapsed amplitude rather than on the limit cycle.

b	τ (s)	$\delta\omega \cdot \tau$	$\omega \cdot \tau$	$\Omega \cdot \tau$	J^*
0.0	4.68	0°	268°	268°	19.93
0.2	4.57	5°	262°	257°	18.10
0.4	4.65	22°	266°	244°	15.48
0.6	5.03	58°	288°	230°	14.16
0.8	6.00	138°	344°	206°	11.53
0.9	(no collapse, $r_{1,\min} = 0.76$)				19.76

7.7 Quantitative Verification: Free Phase Budget

Table 7.2 quantifies the phase accumulation during the collapse interval (defined as the time oscillator 1 spends at $r_1 < 0.5$) at $c = 0.30$ across nonisochronicity values. The free phase $\delta\omega \cdot \tau$ matches the analytical prediction $(\omega - \Omega)\tau$ to within 1° at all b values, confirming the phase velocity surplus mechanism.

The optimizer extends the hold time from 4.7 s ($b = 0$) to 6.0 s ($b = 0.8$) as the per-unit-time phase return increases. The 42% cost reduction from $b = 0$ to $b = 0.8$ is accounted for by 138° of free phase repositioning and the consequent reduction in residual steering cost that must be paid on the limit cycle after amplitude rebuilds.

7.8 Comparison with Direct Optimization

To confirm that the collapse strategy occupies a genuinely distinct cost basin, a structural feature of the control landscape rather than a numerical curiosity, we applied both gradient-based and global stochastic optimization to the same problem at $c = 0.30$ (Table 7.3). We used L-BFGS-B [37], a gradient-based quasi-Newton method with bound constraints, and CMA-ES [38], a derivative-free evolutionary strategy that adapts a covariance matrix over a population of candidate solutions. The control signal was

Table 7.3: Direct optimization vs. HJB at $c = 0.30$.

Method	J^*	$r_{1,\min}$	Evals	Strategy
HJB ($G=64$)	11.53	0.110	—	collapse
L-BFGS-B (zero init)	21.23	0.746	—	on-cycle
L-BFGS-B (warm start)	15.75	0.715	—	on-cycle
L-BFGS-B (5 random)	27.50	0.695	—	on-cycle
CMA-ES (400D)	24.24	0.964	25k	on-cycle
CMA-ES (100D)	24.75	0.991	196k	on-cycle

parameterized as piecewise constant over 100 uniform time intervals.

Neither method discovers the collapse strategy. All L-BFGS-B runs, from zero initialization, a warm start from an adjacent coupling value, and five random initializations, converge to on-cycle solutions with $r_{1,\min} > 0.69$, costing at least 37% more than HJB. CMA-ES was run with 196,000 function evaluations and converged to $r_{1,\min} = 0.99$ at $2.1\times$ the HJB cost. Full-state dynamic programming exposes a lower-cost basin separated from near-cycle local optima that standard direct search does not reach in practice.

7.9 Comparison with Phase-Reduced Control

Phase reduction is the standard approach for oscillator control [8, 7]. To quantify the limitations of reduced-order models, we compare the HJB solution against two phase-reduced baselines.

The *uncoupled baseline* [8] designs controls independently for each oscillator using the infinitesimal phase response curve (iPRC), obtained from the adjoint equation [27, 7]. The *coupled baseline* incorporates the interaction function

$$H(\psi) = \frac{c}{T} \int_0^T Z_x(\phi) [x_\gamma(\phi + \psi) - x_\gamma(\phi)] d\phi$$

Table 7.4: Three-way comparison: uncoupled PRC, coupled PRC, and full-state HJB.

c	$\Delta\phi(T_f)$			J_u		
	Uncoup.	Coupled	HJB	Uncoup.	Coupled	HJB
0.01	178.3°	175.2°	179.7°	7.47	8.14	7.02
0.05	170.8°	177.6°	179.9°	7.47	9.28	8.11
0.10	140.3°	178.6°	180.0°	7.47	10.95	9.63
0.20	48.7°	171.5°	179.9°	7.47	15.07	11.88
0.30	23.7°	171.4°	180.0°	7.47	20.01	11.53
0.50	10.9°	150.9°	180.0°	7.47	30.54	13.07

from [25] and solves the resulting two-dimensional optimal control problem via Pontryagin shooting. Both baselines are applied open-loop to the full four-dimensional system.

Table 7.4 summarizes the comparison. At weak coupling ($c = 0.01$), all three methods perform comparably: the phase difference reaches 175 to 180° at similar cost, and the HJB trajectory stays near the limit cycle. The coupled phase model improves significantly over the uncoupled baseline at moderate coupling: at $c = 0.10$, coupled PRC achieves 178.6° versus 140.3° uncoupled. However, the coupled model saturates near 171° for $c \geq 0.20$ and cannot reach the target, while costing 1.7 to 2.3× more energy than HJB.

The failure of phase reduction here is not merely quantitative but mechanism-changing: the globally optimal strategy at strong coupling does not resemble a better version of phase steering. It uses the full state geometry of the oscillator in a way that no phase-based model can represent. Phase-amplitude reductions [12, 13, 30, 26, 35] recover transverse dynamics, but the collapse trajectory ($r_{1,\min} = 0.11$, $|r - 1| = 0.89$) lies far outside the small-excursion regime in which these methods have been applied and validated. No reviewed reduced-order control study appears to discover or certify a strategy involving excursions of this depth.

Table 7.5: Van der Pol coupling sweep ($\mu = 1$, $u_{\max} = 0.5$, $G = 64$). Off-cycle deviation is measured as $\max_t \min_s \|z_i(t) - x_\gamma(s)\|$, an absolute distance in the (x, y) plane.

c	J^*	Off-cycle 1	Off-cycle 2	$\Delta\phi$
0.05	25.27	0.52	0.51	180.0°
0.10	12.09	0.39	0.39	179.9°
0.15	18.77	0.41	0.58	179.9°
0.20	27.96	0.62	0.80	179.5°
0.25	37.76	1.09	0.86	180.0°
0.30	48.18	1.20	0.91	179.5°

7.10 Validation on Coupled Van der Pol Oscillators

To verify that off-cycle optimal control is not specific to the Stuart-Landau-like algebraic structure, we solved the 4D HJB equation for two coupled Van der Pol oscillators:

$$\dot{x}_i = y_i + c(x_j - x_i) + u_i, \quad \dot{y}_i = \mu(1 - x_i^2)y_i - x_i, \quad (7.11)$$

with $\mu = 1$. The Van der Pol limit cycle is non-circular (egg-shaped), reaching approximately ± 2 in x and ± 2.7 in y , so its “effective amplitude” is roughly twice the unit-circle radius of Stuart-Landau. The natural period at $\mu = 1$ is $T \approx 6.66$ time units. We used $u_{\max} = 0.5$ to match the per-amplitude control authority of the Stuart-Landau setting; with these conventions the off-cycle deviations reported below are absolute deviations from the limit cycle, not deviations normalized by the cycle amplitude.

The Van der Pol oscillator is isochronous in the sense that its angular frequency does not depend on amplitude: near the origin, the linearized dynamics give $\dot{\phi} = 1$, matching the on-cycle average frequency, so $\delta\omega = 0$. This makes Van der Pol a controlled falsification test for the mechanism: the oscillator shares Property 1 (diffusive coupling vanishing at the origin) with Stuart-Landau but lacks Property 2 (phase velocity surplus).

For weak to moderate coupling ($c \leq 0.20$), both oscillators remain close to the limit

cycle, with off-cycle deviations below 0.80 (less than 40% of the limit-cycle amplitude). At $c = 0.25$, the optimal trajectory deviates by 1.09 from the limit cycle (55% of the limit-cycle amplitude), increasing to 1.20 at $c = 0.30$. These significant off-cycle excursions confirm that the HJB controller exploits the full state space even for oscillators with qualitatively different limit-cycle geometry. The cost J^* increases monotonically with c , consistent with $\delta\omega = 0$: there is no free phase accumulation from off-cycle excursions, so no cost non-monotonicity is observed. This confirms both that exploitation of off-cycle dynamics by the HJB controller extends beyond the Stuart-Landau-like algebraic structure, and that the cost non-monotonicity observed in the Stuart-Landau-like case is specifically linked to nonisochronicity.

7.11 Summary

The deterministic results of this chapter establish four facts:

1. There exists a sharp critical coupling c^* above which the optimal control mechanism qualitatively changes.
2. The new mechanism, amplitude collapse, exploits a phase-velocity surplus available near the origin where coupling and nonisochronous distortion both vanish.
3. This mechanism is missed by gradient-based optimization, by global stochastic search, and by phase-reduced control.
4. The mechanism generalizes beyond the Stuart-Landau algebraic structure.

The amplitude-collapse mechanism has two geometric prerequisites. First, *coupling elimination*: the system must possess a state where the synchronizing “torque” vanishes. For diffusive coupling proportional to $c(x_j - x_i)$, this occurs at $r_i \approx 0$. Second, *phase*

velocity surplus: the phase velocity at the origin must exceed the on-cycle average ($\delta\omega = \omega - \Omega > 0$). Condition 1 alone enables collapse (as confirmed by Table 7.1 at $b = 0$ and by the Van der Pol results); Condition 2 determines whether collapse provides an energetic advantage. These two prerequisites distill the mechanism into a form that can be sought in related coupled-oscillator systems beyond the two studied here, and form the basis of the closing discussion in Chapter 8.




Figure 7.2: HJB-controlled trajectories visualized as a three-panel dashboard at $c = 0.10$ (top), $c = 0.20$ (middle), and $c = 0.30$ (bottom). Each row pairs the phase-plane trajectory (left) with the optimal control signals (right). At $c = 0.30$, oscillator 1 spirals inward to a small neighborhood of the origin while oscillator 2 stays near the limit cycle, and the control signal exhibits the three-phase saturated structure described in Section 7.3.1. The transition between the on-cycle regime ($c = 0.10, 0.20$) and the collapse regime ($c = 0.30$) is qualitative rather than quantitative; no intermediate regime exists.




Figure 7.3: Amplitude time histories $r_1(t)$ and $r_2(t)$ at three coupling strengths. At weak coupling, both amplitudes remain near unity. At strong coupling, $r_1(t)$ collapses to ~ 0.1 , holds for a finite window, then rebuilds to the limit cycle while $r_2(t)$ remains near unity. The rebuild rate is consistent with the Floquet exponent $\lambda = -2$ per time unit (Section 7.1.2).




Figure 7.4: Optimal cost $J^*(c)$ across the 30-point coupling sweep. The cost is *non-monotonic* in coupling strength: it peaks near $c \approx 0.22$ and decreases for larger c , despite the stronger coupling opposition. This signature identifies the collapse strategy as energetically advantageous above the critical coupling $c^* = 0.227$.




Figure 7.5: Comparison of phase-reduced and full-state HJB optimal trajectories at three coupling strengths $c = 0.01$ (weak), $c = 0.10$ (moderate), and $c = 0.30$ (strong). At weak coupling, both methods produce similar near-limit-cycle trajectories. At strong coupling, the HJB trajectory leaves the limit-cycle neighborhood entirely while the phase-reduced control cannot represent off-cycle states.

Chapter 8

Conclusions and Future Directions

8.1 Summary of Contributions

This thesis examined a class of optimal control problems for coupled limit-cycle oscillators in which standard methodological assumptions fail simultaneously and in correlated ways. The principal contributions are summarized below.

Deterministic mechanism (Chapter 7). For two diffusively coupled Stuart-Landau-like oscillators, we solved the four-dimensional Hamilton-Jacobi-Bellman equation and identified a critical coupling $c^* = 0.227 \pm 0.003$ above which the globally optimal control mechanism transiently collapses one oscillator’s amplitude to a small neighborhood of the origin. The mechanism exploits a phase-velocity surplus $\delta\omega = b^2/(\omega + \Omega)$ available where the coupling force and the nonisochronous distortion both vanish, and the analytical prediction matches the measured free-phase budget to within 1° across a sweep over nonisochronicity.

Multi-method failure characterization (Chapter 7). The collapse mechanism is missed by gradient-based optimization (L-BFGS-B), by global stochastic search (CMA-ES with up to 1.96×10^5 function evaluations), and by phase-reduced control with or without

an interaction function. The same critical coupling c^* marks the failure boundary for all three families of methods.

Falsification on a non-Stuart-Landau system (Chapter 7). We verified that off-cycle optimal control extends to coupled Van der Pol oscillators, a classic non-Stuart-Landau limit-cycle family with an egg-shaped (non-circular) limit cycle and zero non-isochronicity. The off-cycle excursion persists, confirming that coupling elimination at the origin (Property 1) is sufficient to enable off-cycle control. The cost non-monotonicity disappears, confirming that phase-velocity surplus (Property 2) is responsible for the energetic advantage. Together these two findings decompose the mechanism into two logically independent geometric prerequisites.

Numerical infrastructure (Chapters 5, 6). We extended the CASL-HJX framework [14] to four spatial dimensions with GPU acceleration, achieving HJB solutions in approximately five minutes on a single NVIDIA A100. The implementation is verified by three convergence tests with distinct roles: linear advection establishes fifth-order accuracy of the WENO5 reconstruction; the LQR problem provides an end-to-end pipeline check against the analytical Riccati solution (noting that LQR is not by itself a fifth-order test, since the quadratic value function is captured exactly by any high-order scheme); and Stuart-Landau grid convergence confirms convergence on the actual problem.

8.2 Implications

The results argue that the conventional pipeline of optimal control for nonlinear systems, namely *choose a reduced model, optimize within it, deploy*, can systematically miss the global optimum when the reduced model’s structural assumptions fail. More specifically, the failure is not merely quantitative (the reduced-model optimum is a suboptimal approximation of the true optimum) but qualitative (the optimal mechanism

lies in a region of state space the reduced model cannot represent).

The proposed remedy is straightforward: when computational resources permit, full-state methods such as HJB should be used to verify that the reduced-model optimum is indeed near-optimal. For coupled oscillator control under strong coupling, full-state HJB is the only reference point we know of against which reduced-order and direct-optimization results can be calibrated.

8.3 Limitations

The thesis has several limitations that future work should address.

Two-oscillator scope. The full-state HJB results are restricted to two coupled oscillators because of the curse of dimensionality inherent in grid-based dynamic programming. Whether the collapse mechanism generalizes to three or more oscillators is an open question.

Specific oscillator families. The deterministic results are presented for Stuart-Landau-like oscillators with a specific nonisochronicity structure, with Van der Pol oscillators serving as a controlled cross-check. A systematic study across broader oscillator families (Hodgkin-Huxley, FitzHugh-Nagumo, electrochemical oscillators) has not been performed.

Deterministic-only treatment. The thesis treats the deterministic setting exclusively. Whether and where the amplitude-collapse mechanism survives the addition of intrinsic or extrinsic noise is a question we identify as a natural follow-up but do not address here.

No experimental validation. The mechanism is predicted by the HJB solution but has not been verified in physical hardware. Coupled electrochemical oscillators, coupled lasers, or in-vitro neural preparations would be appropriate platforms.

8.4 Future Directions

Several directions naturally extend the thesis contributions.

Stochastic extension. The deterministic phase-velocity-surplus argument is a statement about expected behavior; whether the surplus accrues on any individual stochastic trajectory is a separate per-realization question. Extending the four-dimensional HJB solver to the stochastic setting using the IMEX-ADI diffusion treatment from our two-dimensional work [24], scaled to 4D via batched cuSPARSE tridiagonal solves on GPU, is in progress.

Calibrated method selection. Reduced-order methods for oscillator control form a hierarchy of increasing fidelity (phase reduction with iPRC, phase reduction with interaction function, phase-amplitude reduction, full-state HJB). A finite-sample distribution-free framework such as conformalized quantile regression [48] could provide calibrated coverage on the per-instance off-cycle deviation and, through that, a certified method-selection decision. This would convert the qualitative observation that “phase reduction fails at strong coupling” into a per-instance diagnostic.

Higher-dimensional oscillator networks. Extending the analysis to $N \geq 3$ coupled oscillators is the most natural follow-up. Methods that may scale beyond the four-dimensional limit of grid-based HJB include tensor-decomposition approaches [41], sparse grids, and neural network parameterizations of the value function [42, 43]. A central question is whether the collapse mechanism generalizes to multiple-oscillator collapse patterns and, if so, how the pattern depends on network topology and target configuration.

Experimental validation. The mechanism predicted by HJB has not been experimentally validated. Suitable platforms include coupled electrochemical oscillators [9], coupled lasers, and in-vitro neural preparations. A successful experimental demonstration

of amplitude collapse would convert the result from a computational prediction into a verified physical phenomenon.

Clinical translation. The Parkinson’s disease motivation that opens this thesis remains. Whether the amplitude-collapse mechanism translates into a clinically realizable DBS protocol is a long-term question requiring close collaboration with neuroscientists and bioengineers. Even an indirect benefit, such as a reduction in DBS duty cycle by exploiting transient amplitude modulation, could meaningfully extend implanted battery life.

Theoretical characterization of the basin separation. The fact that gradient methods and global stochastic search both fail to find the collapse basin suggests an underlying landscape geometry that has not been theoretically characterized. Understanding this geometry, perhaps through Morse-theoretic or topological analysis of the value function, would provide a more rigorous foundation for the empirical observations of this thesis.

8.5 Closing Remarks

The optimal control of coupled oscillators is a well-studied problem with a substantial body of theoretical and computational results. The findings of this thesis suggest that even in this well-studied setting, the conventional methodological pipeline can miss the globally optimal mechanism in a systematic and parameter-dependent way. Identifying where this happens, characterizing the missed mechanism, and verifying it against a non-Stuart-Landau control case are the contributions presented here. Extending the analysis to noisy, higher-dimensional, and physical settings is the work that remains.

Bibliography

- [1] P. A. Tass, *Phase Resetting in Medicine and Biology*. Springer, 1999.
- [2] D. T. Brocker *et. al.*, *Optimized temporal pattern of brain stimulation designed by computational evolution*, *Science Translational Medicine* **9** (2017) eaah3532.
- [3] B. B. Johnson, S. V. Dhople, A. O. Hamadeh, and P. T. Krein, *Synchronization of nonlinear oscillators in an LTI electrical power network*, *IEEE Trans. Circuits Syst. I* **61** (2014), no. 3 834–844.
- [4] G.-S. Seo, M. Colombino, I. Subotic, B. Johnson, D. Gross, and F. Dorfler, *Dispatchable virtual oscillator control for decentralized inverter-dominated power systems*, *IEEE Trans. Power Electron.* **34** (2019), no. 6 5749–5763.
- [5] F. Dorfler and F. Bullo, *Synchronization and transient stability in power networks and nonuniform Kuramoto oscillators*, *SIAM J. Control Optim.* **50** (2012), no. 3 1616–1642.
- [6] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence*. Springer, 1984.
- [7] B. Monga, D. Wilson, T. Matchen, and J. Moehlis, *Phase reduction and phase-based optimal control for biological systems: A tutorial*, *Biol. Cybern.* **113** (2019) 11–46.
- [8] J. Moehlis, E. Shea-Brown, and H. Rabitz, *Optimal inputs for phase models of spiking neurons*, *ASME J. Comput. Nonlinear Dyn.* **1** (2006) 358–367.
- [9] A. Zlotnik, Y. Chen, I. Z. Kiss, H. A. Tanaka, and J.-S. Li, *Optimal waveform for fast entrainment of weakly forced nonlinear oscillators*, *Phys. Rev. Lett.* **111** (2013) 024102.
- [10] D. Wilson and J. Moehlis, *Isostable reduction of periodic orbits*, *Phys. Rev. E* **94** (2016) 052213.
- [11] D. Wilson, *Phase-amplitude reduction far beyond the weakly perturbed paradigm*, *Phys. Rev. E* **101** (2020) 022220.
- [12] D. Wilson, *An adaptive phase-amplitude reduction framework without $O(\epsilon)$ constraints on inputs*, *SIAM J. Appl. Dyn. Syst.* **19** (2020) 204–230.

- [13] A. N. Dewanjee and D. Wilson, *Optimal phase-based control of strongly perturbed limit cycle oscillators using phase reduction techniques*, *Phys. Rev. E* **109** (2024) 024223.
- [14] F. Rajabi, J. Fingerman, A. Wang, J. Moehlis, and F. Gibou, *CASL-HJX: A comprehensive guide to solving deterministic and stochastic Hamilton-Jacobi equations*, *Comput. Phys. Commun.* **319** (2026) 109941.
- [15] F. Rajabi, F. Gibou, and J. Moehlis, *Beyond phase reduction: Amplitude collapse in optimal control of coupled oscillators*, in *Submitted to IEEE Conf. Decision and Control (CDC)*, 2026.
- [16] S. Little, A. Pogosyan, S. Neal, B. Zavala, L. Zrinzo, M. Hariz, *et. al.*, *Adaptive deep brain stimulation in advanced Parkinson disease*, *Ann. Neurol.* **74** (2013), no. 3 449–457.
- [17] A. Oswal, P. Brown, and V. Litvak, *Synchronized neural oscillations and the pathophysiology of Parkinson’s disease*, *Curr. Opin. Neurol.* **26** (2013), no. 6 662–670.
- [18] J. M. Hurtado, L. L. Rubchinsky, K. A. Sigvardt, V. L. Wheelock, and C. T. E. Pappas, *Temporal evolution of oscillations and synchrony in GPi/muscle pairs in Parkinson’s disease*, *J. Neurophysiol.* **93** (2005), no. 3 1569–1584.
- [19] S. F. Lempka, M. D. Johnson, S. Miocinovic, J. L. Vitek, and C. C. McIntyre, *Current-controlled deep brain stimulation reduces in vivo voltage fluctuations observed during voltage-controlled stimulation*, *Clin. Neurophysiol.* **121** (2010), no. 12 2128–2133.
- [20] M. Arlotti, M. Rosa, S. Marceglia, S. Barbieri, and A. Priori, *The adaptive deep brain stimulation challenge*, *Parkinsonism Relat. Disord.* **28** (2018) 12–17.
- [21] A. Priori, G. Foffani, L. Rossi, and S. Marceglia, *Adaptive deep brain stimulation (aDBS) controlled by local field potential oscillations*, *Exp. Neurol.* **245** (2013) 77–86.
- [22] A. Nabi, M. Mirzadeh, F. Gibou, and J. Moehlis, *Minimum energy desynchronizing control for coupled neurons*, *J. Comput. Neurosci.* **34** (2013) 259–271.
- [23] D. Wilson and J. Moehlis, *A Hamilton-Jacobi-Bellman approach for termination of seizure-like bursting*, *J. Comput. Neurosci.* **37** (2014) 345–355.
- [24] F. Rajabi, F. Gibou, and J. Moehlis, *Optimal control for stochastic neural oscillators*, *Biol. Cybern.* **119** (2025), no. 2 Article 9.
- [25] B. Ermentrout and D. Terman, *Mathematical Foundations of Neuroscience*. Springer, 2010.

- [26] B. Monga and J. Moehlis, *Optimal phase control of biological oscillators using augmented phase reduction*, *Biol. Cybern.* **113** (2019) 161–178.
- [27] E. Brown, J. Moehlis, and P. Holmes, *On the phase reduction and response dynamics of neural oscillator populations*, *Neural Comput.* **16** (2004) 673–715.
- [28] D. Wilson and J. Moehlis, *Optimal chaotic desynchronization for neural populations*, *SIAM J. Appl. Dyn. Syst.* **13** (2014) 276–305.
- [29] D. Wilson and J. Moehlis, *Recent advances in the analysis and control of large populations of neural oscillators*, *Ann. Rev. Control* **54** (2022) 327–351.
- [30] D. Wilson and B. Ermentrout, *Augmented phase reduction of (not so) weakly perturbed coupled oscillators*, *SIAM Rev.* **61** (2019), no. 2 277–315.
- [31] A. Franci, W. Pasillas-Lepine, and A. Chaillet, *Validity of the phase approximation for coupled nonlinear oscillators: A case study*, in *Proc. IEEE Conf. Decision Control*, pp. 2316–2321, 2012.
- [32] D. Wilson, *Optimal control of oscillation timing and entrainment using large magnitude inputs: An adaptive phase-amplitude-coordinate-based approach*, *SIAM J. Appl. Dyn. Syst.* **20** (2021) 1814–1843.
- [33] N. Namura and H. Nakao, *Optimal phase control of limit-cycle oscillators with strong inputs through phase-amplitude reduction*, in *Proc. IEEE 63rd Conf. Decision and Control (CDC)*, pp. 4003–4009, 2024.
- [34] R. Nicks, R. Allen, and S. Coombes, *Insights into oscillator network dynamics using a phase-isostable framework*, *Chaos* **34** (2024) 013141.
- [35] K. Toth and D. Wilson, *Direct method phase-based model identification for coupled oscillators*, *Physica D* **484** (2025) 134968.
- [36] A. Pérez-Cervera, B. Gutkin, P. J. Thomas, and B. Lindner, *A universal description of stochastic oscillators*, *Proc. Natl. Acad. Sci. USA* **120** (2023), no. 29 e2303222120.
- [37] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, *A limited memory algorithm for bound constrained optimization*, *SIAM J. Sci. Comput.* **16** (1995), no. 5 1190–1208.
- [38] N. Hansen, *The CMA evolution strategy: A tutorial*, *arXiv:1604.00772* (2016).
- [39] D. E. Kirk, *Optimal Control Theory: An Introduction*. Dover Publications, 1970.
- [40] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. 1. Athena Scientific, 2nd ed., 1995.

- [41] S. Dolgov, D. Kalise, and K. Kunisch, *Tensor decomposition methods for high-dimensional Hamilton-Jacobi-Bellman equations*, *SIAM J. Sci. Comput.* **43** (2021) A1625–A1650.
- [42] T. Nakamura-Zimmerer, Q. Gong, and W. Kang, *Adaptive deep learning for high dimensional Hamilton-Jacobi-Bellman equations*, *arXiv:1907.05317* (2019).
- [43] J. Han, A. Jentzen, and W. E, *Solving high-dimensional partial differential equations using deep learning*, *Proc. Natl. Acad. Sci. USA* **115** (2018), no. 34 8505–8510.
- [44] S. Osher and R. Fedkiw, *Level Set Methods and Dynamic Implicit Surfaces*, vol. 153. Springer, 2003.
- [45] G. Barles and P. E. Souganidis, *Convergence of approximation schemes for fully nonlinear second order equations*, *Asymptotic Analysis* **4** (1991), no. 3 271–283.
- [46] C.-W. Shu and S. Osher, *Efficient implementation of essentially non-oscillatory shock-capturing schemes*, *J. Comput. Phys.* **77** (1988) 439–471.
- [47] M. Golubitsky, I. Stewart, and D. G. Schaeffer, *Singularities and Groups in Bifurcation Theory*, vol. 2. Springer, 1988.
- [48] Y. Romano, E. Patterson, and E. J. Candès, *Conformalized quantile regression*, in *Advances in Neural Information Processing Systems*, 2019.