

Abstract

Dynamical Behavior near Exceptional Points in an Optomechanical System

David R. Mason

2018

Coupled mechanical oscillators have long been an archetypical system for understanding eigenmodes and coupled dynamics. But in the last few decades, the study of open systems (i.e. those open to loss or gain) has brought a fresh interest and perspective to such simple systems, revealing a surprisingly rich set of physical phenomena. Specifically, it was realized that degeneracies in open systems ('exceptional points', or EPs) possess a non-trivial topology, with interesting implications for closed adiabatic cycles. The theoretical properties of EPs have been made increasingly clear over the last 20 years, but experimental progress has generally been limited to spectroscopy, with no demonstrations of the predicted dynamical behavior. Here, I'll present work in which we use a cavity optomechanical system as a convenient, highly tunable platform for studying this multimode physics.

I'll begin with a pedagogical introduction to cavity optomechanics, including our particular experimental realization: a Si_3N_4 membrane coupled to a high-finesse optical cavity. Then, the physics of exceptional points will be reviewed using a toy model, before seeing how these features are accessible in our optomechanical system. I'll then present our study of time-dependent perturbations of this system, which provided the first experimental demonstration of adiabatic (and non-adiabatic) behavior near an EP. These perturbations can be used to affect energy transfer which is both topology-dependent and non-reciprocal. This demonstration relies on a somewhat fortunate symmetry in our system, but in the final chapter, we'll see that this restriction can be lifted, to enable this energy transfer in a broad class of systems.

Dynamical Behavior near Exceptional Points in an Optomechanical System

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
David R. Mason

Dissertation Director: Jack Harris

May, 2018

Copyright © 2017 by David R. Mason
All rights reserved.

Contents

Contents	iii
List of Figures	vi
Acknowledgements	ix
1 Optomechanics	1
1.1 Introduction	1
1.2 Optical cavities	3
1.2.1 Overview of optical resonators	3
1.2.2 Mathematical model for an optical cavity	4
1.3 Mechanical oscillators	7
1.3.1 Overview and context	7
1.3.2 Mathematical model for a mechanical oscillator	7
1.3.3 Rewriting as a complex amplitude	9
1.4 Optomechanical Coupling	11
1.4.1 Canonical system	11
1.4.2 Solving the optomechanical system: Linearization	14
1.4.3 Solving the optomechanical system: Analysis in Fourier domain . . .	15
1.4.4 Optomechanical Control: Dynamical Backaction via Σ	18
1.5 Membrane-in-the-Middle Optomechanics	21
1.5.1 Model: Dielectric slab within cavity	22
1.5.2 Implementation: Si ₃ N ₄ Membranes	24

1.5.3	Historical context: Membranes at Yale	26
1.5.4	State-of-the-field: Membranes elsewhere	27
1.5.5	Previous Optomechanics Work with this System	29
2	Experimental Setup	32
2.1	Cryogenic Optomechanical Platform	32
2.1.1	Cryostat	33
2.1.2	Optical Cavity	34
2.1.3	Membrane	35
2.1.4	Inside the Cryostat	36
2.2	Optical Setup	39
2.3	Detection Electronics	42
2.4	Feedback Techniques	44
3	Exceptional Points: Degeneracies in Open Systems	48
3.1	Closed Systems: Avoided Crossings and the Adiabatic Theorem	49
3.2	Open Systems & Complex Eigenvalues	52
3.2.1	Coupled Modes in an Open System	52
3.2.2	Alternative Perspective: Complex Avoided Crossings	55
3.3	Adiabatic Control in Open Systems	56
3.3.1	Dependence on Eigenvalue Topology	60
3.3.2	Dependence on Perturbation Speed	61
3.4	Non-adiabatic behavior	63
3.5	Context and State-of-the-Field	68
4	Optomechanical Exceptional Points	74
4.1	Optomechanics with Multiple Mechanical Modes	75
4.1.1	Motivation for coupled-mode model	75
4.1.2	Deriving mode coupling	76
4.1.3	Mechanical Eigenmodes: Optical Spring, Damping, and Coupling . .	78
4.2	Experimental Measurement of Eigenvalues	80

4.3	Studying Dynamical Behavior: Initialization, Perturbation, and Measurement	84
4.3.1	Details of Energy Transfer Measurement	85
4.4	Analysis: Varying Loop Speed and Initial Condition	90
4.5	Analysis: Dependence on Loop Geometry	93
5	Virtual Exceptional Points from Non-Degenerate Modes	95
5.1	Coupling through a Bichromatic Drive: Motivation	96
5.2	Coupling through a Bichromatic Drive: Derivation	97
5.3	Optomechanical eigenvalues in presence of bichromatic drive	106
5.4	Experimental Implementation	109
5.5	Measuring Mechanical Eigenvalues	110
5.5.1	Fitting Overlapping Modes	111
5.6	Adiabatic Energy Transfer near a VEP	116
6	Conclusion and Outlook	119
A	Dominance of Gain Mode Beyond Adiabatic Regime	122
	Bibliography	125

List of Figures

1.1	Optical Resonators: Overview	3
1.2	Optical Resonances: Coupling Regimes	6
1.3	Mechanical Resonances: Overview	9
1.4	Canonical Optomechanical System	12
1.5	Dynamical Backaction in Different Sideband Regimes	19
1.6	Dynamical Backaction: Spanning the Complex Plane	21
1.7	Membrane-in-the-Middle Optomechanics	22
1.8	Solving for Dielectric Perturbation of a Fabry-Perot Cavity	23
1.9	Membrane Vibrational Modes	26
1.10	Measured Thermal Motion Spectrum	26
1.11	Si ₃ N ₄ Membrane Photos	28
1.12	Prior Work: Optical Avoided Crossings	29
1.13	Prior Work: Observation of Sideband Asymmetry	31
2.1	³ He Cryostat Schematic	33
2.2	Cryogenic Platform Schematic	34
2.3	Mechanical Ringdown Measurements	36
2.4	Photos of Experimental Setup	37
2.5	Membrane Position Dependence of Cavity Modes	39
2.6	Optical Circuit and Optical Frequency Layout	40
2.7	Filter Cavity Design	41
2.8	Electrical Measurement Circuit	43

2.9	PDH Error Signal	45
2.10	Feedback Circuit Schematic	46
3.1	Closed System Eigenvalues	50
3.2	Open System Eigenvalues: λ_{\pm} vs. Δ	53
3.3	Open System Eigenvalue Sheets	54
3.4	Monodromy: Example Trajectories on Eigenvalue Sheets	54
3.5	Eigenvalues in the Complex Plane	55
3.6	Trajectory Simulation: Adiabatic Energy Transfer	59
3.7	Energy Transfer Dependence on Loop Geometry	60
3.8	Simulated Trajectory for Diabatic Perturbation	62
3.9	Energy Transfer Dependence on Loop Speed	62
3.10	Calculated Trajectory with Non-Adiabatic Jump	64
3.11	Calculated Propagation Matrix Elements	66
3.12	Simulated Clockwise Trajectories: Adiabatic and Non-Adiabatic Behavior	67
3.13	Energy Transfer Asymmetry: Summary	68
3.14	Summary of Early Experimental Results	70
4.1	Nearly-Degenerate $\{1, 3\}/\{3, 1\}$ Modes	75
4.2	Optomechanical Eigenvalue Sheets	78
4.3	Calculated Optomechanical Eigenvalues in the Complex Plane	79
4.4	Calculated Optomechanical Eigenvalues in the Complex Plane: Zoom-In	80
4.5	Driven Response Measurements for Extracting Mechanical Eigenvalues	81
4.6	Measured Mechanical Eigenvalues in the Complex Plane	82
4.7	Measured and Calculated Eigenvalue Sheets	83
4.8	Illustration of Energy Transfer Measurement	84
4.9	Energy Transfer Measurement Protocol	85
4.10	Electrical Circuit for Energy Transfer Measurements	87
4.11	Measurement Averaging Procedure	88
4.12	Analysis: Quantifying Energy Transfer	89

4.13 Measured Energy Transfer Dependence on Adiabaticity, Initial Condition, and Loop Direction	91
4.14 Measured Propagator Matrix Element Amplitudes	93
4.15 Measured Energy Transfer Dependence on Loop Geometry	94
5.1 Bichromatic Drive Scheme: Frequency Layout	96
5.2 Calculated Eigenvalue Sheets in Degenerate Frame	109
5.3 Electrical Circuit Schematic for Accessing VEP	110
5.4 Driven Response Measurements of Non-Degenerate Modes	111
5.5 Hybridized Motion in Driven Response Measurements	112
5.6 Driven Response Measurements and Fits	113
5.7 Measured Eigenvalue Sheets near VEP	114
5.8 Measured Eigenvalue Differences near VEP	115
5.9 Adiabatic Energy Transfer by Encircling VEP	116
5.10 VEP Energy Transfer Dependence on Loop Geometry	117
5.11 Measured VEP Energy Transfer Dependence on Adiabaticity, Initial Condi- tion, and Loop Direction	118
A.1 Trajectory Simulation: Appearance of Additional Non-Adiabatic Jumps . .	123
A.2 Very-Long-Time Energy Transfer Asymmetry: Summary	124

Acknowledgements

This thesis represents the cumulative work of six years, during which time I've been lucky to have the help and support of many people.

I'm very grateful to have had the opportunity to work with and learn from Jack Harris during my time at Yale. Jack brings a careful, thoughtful approach to every part of the scientific process, and I hope to carry this with me throughout my career. He is a meticulous experimentalist, a patient and insightful teacher, and on top of it all, a really nice guy. Thanks for everything, Jack. I'll always look forward to seeing what new ideas come out of the Harris Lab.

I'm also thankful to have worked with my original cryo-optomechanics teammates: Donghun Lee and Woody Underwood. The success of our complex cryogenic apparatus is a direct result of Donghun's careful design, and any understanding I have of electrical circuits is thanks to Woody's help over the years. Moreover, Donghun and Woody are two of the kindest people I've met, and working with them was always a pleasure. I have no doubt that in his new lab, Donghun's patient mentoring will produce a new generation of great scientists, and Woody's technical expertise will help advance optical technology, as well as cabin construction. Thanks as well to the newer members of the membrane experiment, Luyao Jiang and Haitan Xu, who helped carry out the work in this thesis. I wish you both the best of luck exploring new science with the contents of the big blue dewar in SPL17!

I learned from many other Harris Lab members over the years as well. For sure, I owe much of my understanding of optomechanics (and experimental physics in general) to Alexey Shkarin, whose willingness to answer unending questions on the way to the coffee room was nothing short of heroic. His intuition for both experimental and theoretical problems is

remarkable, and whatever field he ends up in will be lucky to have him. Anya Kashkanova was also always willing to think through problems together and offer new perspectives. She is a careful scientist and a patient teacher, and as with Alexey, will I'm sure have success wherever she goes. Countless other members of Harris Lab helped in this same way: Anthony Lollo, Glen Harris, Charles Brown, Kale Johnson, Ivana Petkovic, Nathan Flowers-Jacobs, Dustin Ngo, Scott Hoch – I'm certain that conversations and arguments across the office with these people taught me as much as all the papers I read during grad school.

Beyond physics, those same people also made Harris Lab a wonderful place to work, offering friendly faces to share lunch with every day (an important part of surviving graduate school). In the case of Anya, Alexey, and Anthony, we not only shared lunches, but also breakfasts, dinners, movies, adventures, and a rather absurd number of board games. And while Harris Lab is full of wonderful people, I'm also grateful for the many friends from outside optomechanics who helped maintain sanity over the years: Corey and Shannon Adams, Evan Pease, Ryan Timpe, Carl T. West, Jaymin Patel, Rona Ramos, Whit Schofield, et. al. – Thanks for filling New Haven (and 82 Mechanic Street, in particular) with so many wonderful memories.

I'm also grateful to have shared most of the last 6 years with Hannah Fuson, who has been a constant source of support (even when 186 miles away). New Haven will always be full of many happy (and tasty) memories from our time here. Thanks for all the adventures over these last years, our newest one in Denmark, and the many yet to come.

Finally, I'm grateful to my parents and family. Making the trip back to Frostburg has always offered a warm, refreshing break from life in grad school. I'm thankful to my parents for a childhood full of books, and for tolerating my various moves away from home in pursuit of a somewhat odd-seeming career path. I promise I ~~will~~ may eventually have something that looks like a real job.

Chapter 1

Optomechanics

In principle, the physics presented in this thesis (i.e. the physics of exceptional points) is nothing more than a careful study of parametrically coupled, damped harmonic oscillators. Yet remarkably, this simple system reveals rich phenomena associated with topology, adiabatic control, and complex eigensystems. This unexpected depth from mere coupled oscillators only becomes visible if one has sufficient control over their resonance parameters and can measure the motion with sufficient accuracy. Given these requirements, cavity optomechanical devices offer a perfect platform for exploring these multimode interactions, as we shall see in this chapter.

1.1 Introduction

Optomechanics involves the use of electromagnetic fields (typically lasers or microwave fields) to measure and control the motion of mechanical oscillators [1]. Experimental realizations range from km/kg-scale interferometers for detecting gravitation waves to the motion of trapped atoms, thus spanning many orders of magnitude in scale. Over the last 10 years, rapid technological advances have made it possible for optomechanical systems to begin routinely functioning as unequivocally-quantum systems. The field now frequently sees oscillators cooled to their quantum ground state, with clear non-classical experimental signatures. Furthermore, having reached this regime, various optomechanical systems have now been able to recreate some of the hallmarks of quantum optics, like squeezing and

entanglement. As the field pushes further in this direction (towards non-Gaussian quantum states, like Fock states and superpositions), these devices will be poised to probe new physics associated with the quantum behavior of massive objects.

In parallel with this progress towards creating quantum mechanical oscillators, the field has also sought to exploit the properties of optomechanical systems for other applications, both quantum and classical. One strong theme in this research is to use mechanical oscillators as interfaces: Instead of treating the motion itself as the object of study, one can couple the motion to some other desired degrees of freedom, and thereby enable optical readout/control through the optomechanical interaction. The most straightforward example of this transduction might be the use of optomechanical systems as force sensors, but other work has coupled mechanical motion to magnetic fields, spins, atomic degrees of freedom, and gravitational waves. Mechanical oscillators also present one of the more promising approaches to convert signals between the microwave and optical domains – an important capability for networking superconducting qubit systems.

While working towards these various goals, the field of optomechanics has developed systems which offer the ability to measure and manipulate mechanical oscillators with unprecedented accuracy and control. It is this simple advantage that will be exploited in the work presented in this thesis. To understand these capabilities, we will begin with a pedagogical introduction to the field of optomechanics. We'll first consider separately the basic concepts of optical cavities and mechanical resonators, including both their physical realization and mathematical description. Then, we'll examine the basic model used to describe a coupled optomechanical system, paying attention to how the ideas of measurement and control are manifest in this model. Finally, we'll discuss the particular optomechanical system used in this thesis: a cryogenic membrane-in-the-middle system. We'll outline its theoretical basis, discuss the role such systems play in the larger field, and note the historical context of this particular experiment.

1.2 Optical cavities

1.2.1 Overview of optical resonators

Optical cavities are objects which spatially and temporally confine electromagnetic waves, usually by repeated reflections off a boundary. This confinement allows for stronger interactions with systems placed inside the cavity, such as atoms, non-linear materials, and as we will see later, mechanical oscillators. Moreover, this confinement is a resonant process occurring only for particular frequencies, making optical cavities useful as narrow-bandwidth sources (lasers) and filters.

Traditionally, optical cavities were formed by a combination of highly-reflective mirrors, as illustrated in Fig 1.1. This configuration is referred to as a Fabry-Perot resonator, after Charles Fabry and Alfred Perot – two French physicists who theoretically and experimentally introduced the now-canonical system in 1899 [2]. Two reflective mirrors result in the formation of a standing wave, with the resonance condition requiring that the cavity length be spanned by an integer number of wavelengths. The cavity will support many such resonances, spaced in frequency by the 'free spectral range', $\text{FSR} = \frac{c}{2L}$ where c is the speed of light and L is the distance between the mirrors. The inverse of the FSR corresponds to the time for light to make one round trip inside the cavity.

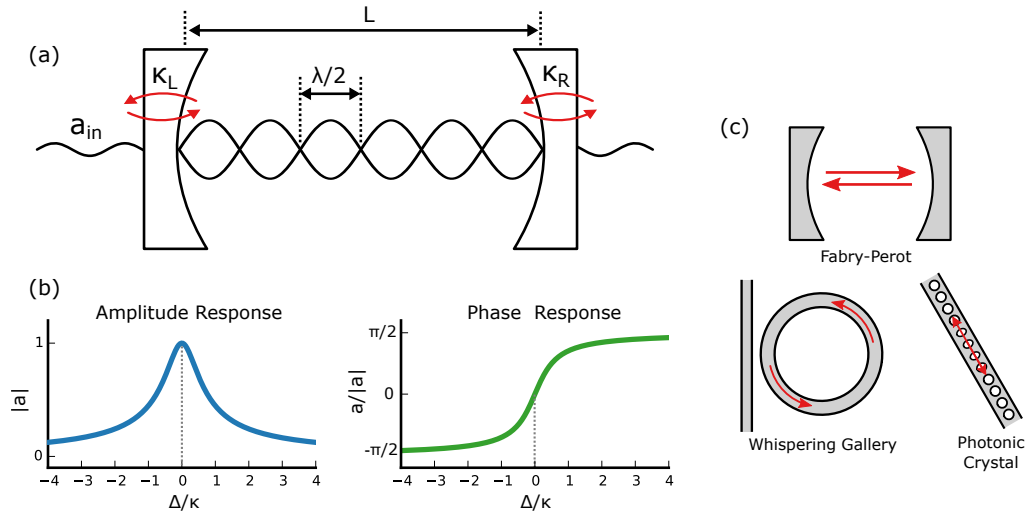


Figure 1.1: Optical Resonators. (a) Schematic of Fabry-Perot cavity (b) Cavity amplitude and phase response. (c) Optical resonator geometries, including Fabry-Perot, a whispering gallery resonator (driven by a tapered waveguide), and a photonic crystal cavity

Due to absorption/loss and the finite transmissivity of the mirrors, one expects that light will remain in the cavity for a finite amount of time. The cavity loss can be characterized in different ways, depending on context and emphasis. If τ_{cav} is the characteristic time for light to decay from the cavity, then one can define the cavity decay rate, $\kappa = 1/\tau_{\text{cav}}$, which is particularly relevant in describing the cavity's spectral response (we will see later that is a Lorentzian resonance of width κ). One can also discuss the finesse ($\mathcal{F} = \frac{\tau_{\text{cav}}}{1/FSR} = \frac{FSR}{\kappa}$) of the cavity, which corresponds to the mean number of reflections the light undergoes while in the cavity. This definition also highlights that \mathcal{F} describes the relative width of a single resonance with respect to the overall spacing of resonances. Finally, one might also define the Q -factor of the cavity, defined as the ratio of the field's oscillation rate to its loss rate: $Q = \frac{\omega_c}{\kappa}$. Using our prior definitions, one sees that this quantity also corresponds to the number of times the light field oscillates before leaving the cavity: $Q = \frac{\tau_{\text{cav}}}{1/\omega_c}$.

Fabry-Perot cavities still form the basis of experiments throughout atomic physics and quantum optics, and modern advances continue to push them towards lower loss and greater stability. At the same time, alternative cavity geometries have also been developed, based on total-internal reflection within circular/spherical dielectric structures. These ‘whispering gallery’ resonators are able to achieve high- Q resonances in a monolithic, chip-based photonic platform. The last few decades have also seen growth in the field of photonic crystal cavities, in which one generates optical confinement in 2D dielectric slabs via the engineering of a quasi-periodic optical potential. Both whispering gallery and photonic crystal cavities offer the ability to achieve high- Q resonances with microscopic mode volumes, which is valuable for some applications.

1.2.2 Mathematical model for an optical cavity

For our purposes, a single mode of the electromagnetic field inside a high- Q optical cavity will be described by a complex amplitude a . Classically, this variable corresponds to the mode's electric field amplitude, and its dynamics can be derived from Maxwell's equations and application of appropriate boundary conditions. Alternatively, in a quantum treatment, one can model the cavity as a bosonic mode coupled to an external bath, tracing over the bath modes and solving for the Heisenberg equation of motion. In the classical limit, these

descriptions are equivalent, and we will simply adopt the final equation of motion as our starting point. The cavity is parameterized by a frequency ω_c and two loss/coupling rates, κ and κ_{in} . The latter (κ_{in}) describes the coupling rate of the input mirror, through which the cavity can be driven and monitored. The total loss rate, κ , includes κ_{in} as well as coupling to the other loss channels (e.g. the output mirror or absorption processes). The dynamics of the cavity are given by:

$$\dot{a} = \left(-\frac{\kappa}{2} - i\omega_c\right)a + \sqrt{\kappa_{\text{in}}}a_{\text{in}} \quad (1.1)$$

Here, a_{in} denotes whatever drive field we apply via the input mirror. If we consider a continuous coherent drive $a_{\text{in}}(t) = \bar{a}_{\text{in}}e^{-i\omega_L t}$ and look for a steady-state solution in which the cavity responds at the drive frequency ($a(t) = \bar{a}e^{-i\omega_L t}$), we find that the field inside the cavity is given by:

$$\bar{a} = \frac{\sqrt{\kappa_{\text{in}}}}{\frac{\kappa}{2} - i\Delta} \bar{a}_{\text{in}} \quad (1.2)$$

where $\Delta = \omega_L - \omega_c$ is defined as the detuning of the drive. Note that standard practice in the field is to normalize the cavity and input fields in terms of their photon number and photon flux, respectively. That is, $|\bar{a}|^2 = N_{\text{cav}}$ and $|\bar{a}_{\text{in}}|^2 = \frac{P_{\text{in}}}{\hbar\omega_L}$ where N_{cav} is the intracavity photon number and P_{in} is the incident power of the drive. This normalization in terms of photons is not necessary, but allows one to easily move to a quantum description, and we will follow it here to conform with standard notation.

The form of \bar{a} already highlights the key features of an optical cavity: Namely, a resonant response as the laser drive frequency is brought near the cavity frequency. The amplitude and phase of this response is sketched in Fig. 1.1. The intracavity power ($\propto |\bar{a}|^2$) traces out a Lorentzian, centered at $\Delta = 0$ and with width κ . The phase also evolves by π as the cavity detuning is swept, with a maximal linear phase change near $\Delta = 0$. This rapid change in amplitude/phase near resonance is what allows cavities to function as sensitive transducers of quantities which couple to the cavity frequency.

In the work presented in this thesis, we will choose to interrogate our optical cavity by measuring its reflected lightF, in which case our signal will contain two components: the

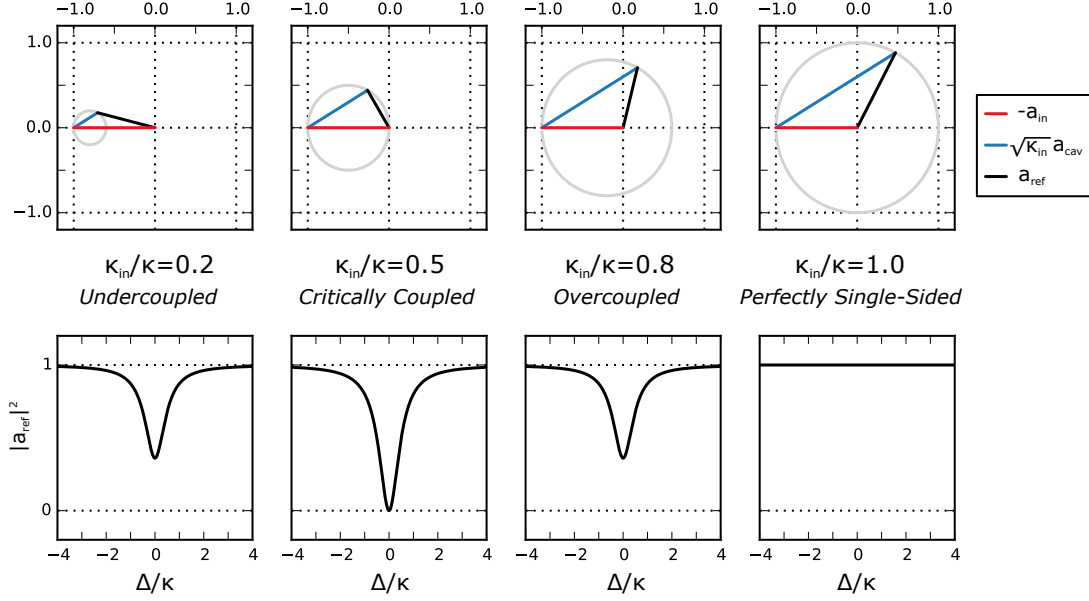


Figure 1.2: Cavity Response in Different Coupling Regimes. Top Row: Complex response of the intracavity field (a_{cav}) and the reflected field (a_{ref}), as detuning is swept from $\Delta \ll 0$ to $\Delta \gg 0$. Bottom Row: Reflected power, $|a_{\text{ref}}|^2$. In all plots, $a_{\text{in}} = 1$ and $\kappa = 1$

promptly reflected cavity drive (which acquires a π phase on reflection), and cavity light leaking out through the input mirror, at a rate κ_{in} . The total reflected beam will be:

$$a_{\text{ref}} = a_{\text{in}} e^{i\pi} + \sqrt{\kappa_{\text{in}}} a \quad (1.3)$$

$$a_{\text{ref}} = -\left(1 - \frac{\kappa_{\text{in}}}{\frac{\kappa}{2} - i\Delta}\right) a_{\text{in}} \quad (1.4)$$

In this expression, we see that the cavity light interferes destructively with the prompt reflection, tracing out a complex trajectory as in Fig. 1.2. If one is only measuring the power of the reflected light, one sees a Lorentzian dip as the laser detuning is swept.

The size of κ_{in} relative to κ determines the extent of the interference, and defines several coupling regimes, summarized in Fig. 1.2. For instance, in an ideal (lossless) symmetric cavity (in which both mirrors are equally reflective), κ_{in} will be make up exactly half of the total loss, leading to perfect destructive interference with the prompt reflection: $\overline{a_{\text{ref}}} = 0$ for $\Delta = 0$. This condition defines a ‘critically coupled’ cavity. If the input mirror is the main source of loss from the cavity ($\kappa_{\text{in}} \rightarrow \kappa$), then on resonance the leaked field will be twice the

magnitude of the input field, leading to $\overline{a_{\text{ref}}} = +\overline{a_{\text{in}}}$. This defines a maximally ‘overcoupled’ (or ‘single-sided’) cavity, in which there is actually no resonant dip in reflection, only a phase shift. This regime is desirable for many quantum protocols, in which it is important to collect as much of the cavity light as possible. Finally, $\kappa_{\text{in}} \ll \kappa$ defines an ‘undercoupled’ cavity. Note in Fig. 1.2 that just measuring reflected power is insufficient to uniquely measure κ_{in} , since undercoupled and overcoupled cavities can produce the same dip. One must additionally measure transmission or reflected phase in order to fully characterize a cavity.

1.3 Mechanical oscillators

1.3.1 Overview and context

The conceptually-simple mechanical oscillator has long played an important role in science and technology. Indeed, the field of time-keeping, as just one example, has been built around mechanical resonators for centuries, dating back to Galileo’s study of the pendulum in 1602. The humble pendulum clock would remain the de-facto standard for precision time-keeping until the 1940s¹, when it was displaced by the Quartz oscillator – yet another mechanical system. This shift to a piezoelectric crystal oscillator might be seen as an early milestone in the development of MEMs (microelectromechanical systems) technologies, which now pervade modern electronics, both for timekeeping and sensing. Optomechanics is, in some ways, an evolution of this quest to expand the capabilities of mechanical systems, pushing toward unprecedented levels of precision and control. Yet underlying all of this physics and engineering is the same simple harmonic oscillator, whose properties we review here.

1.3.2 Mathematical model for a mechanical oscillator

We will begin with the simplest model of a mechanical oscillator, noting some general properties and proceeding to manipulate the model into a form which will be more convenient

1. Here, a historical reminder that high-Q mechanical resonators are not a particularly modern achievement: The Shortt-Synchronome clock, invented in 1921, had a Q -factor of $\sim 100,000$, and stood as the most accurate commercially-produced mechanical clock until the 1940s. It was sufficiently accurate to measure seasonal variations in the Earth’s rotation, as well as variations in gravity due to tidal forces.

later on. We start with a system of mass m , with resonance frequency ω_m . The system is assumed to be subject to some linear damping at rate γ_m , and a driving force F . We note that the damping can also be described by the mechanical Q -factor, defined as the ratio of oscillation rate to energy loss rate: $Q = \omega_m/\gamma_m$. The equation of motion for the system is

$$m\ddot{x} = -m\omega_m^2 x - m\gamma_m \dot{x} + F \quad (1.5)$$

Before going any further, we note the solutions to this, in both the time and frequency domains. By looking for steady-state solutions $e^{\lambda t}$ of the homogeneous equation ($F = 0$), we find that the eigenvalues λ of the system are given by

$$\lambda = \frac{1}{2} \left(\gamma_m \pm \sqrt{\gamma_m^2 - 4\omega_m^2} \right) \quad (1.6)$$

From this, it is clear that for $\gamma_m > 2\omega_m$ yields purely real exponential solutions, while small damping ($\gamma_m < 2\omega_m$) yields complex eigenvalues with oscillating solutions².

Alternatively, we can consider the Fourier transform of Eq. 1.5, to understand the system's response to a drive with spectrum $F[\omega]$:

$$x[\omega] = \frac{F[\omega]/m}{\omega_m^2 - \omega^2 + i\gamma_m\omega} = \chi_x[\omega]F[\omega] \quad (1.7)$$

Here we see the full complex response function of the damped harmonic oscillator, which can be broken up into the amplitude and phase response, plotted in Fig. 1.3. This indicates the amplitude and phase with which the oscillator would respond to a sinusoidal drive.

If the mechanical Q -factor is very large ($\omega_m \gg \gamma_m$, and one is only interested in the response of the system near resonance, then it is often sufficient to work with an approximated form of Eq. 1.7. Specifically, we can approximate $\omega_m^2 - \omega^2 = (\omega_m - \omega)(\omega_m + \omega) \approx (\omega_m - \omega)(2\omega_m)$, so that Eq. 1.7 becomes

$$x[\omega] \approx \frac{F[\omega]/m}{2\omega_m^2 - 2\omega\omega_m + i\gamma_m\omega} = \frac{F[\omega]/(2m\omega_m)}{\omega_m - \omega + i\frac{\gamma_m}{2}\frac{\omega}{\omega_m}} \approx \frac{F[\omega]/(2m\omega_m)}{\omega_m - \omega + i\frac{\gamma_m}{2}} \quad (1.8)$$

2. Of course, $x(t)$ cannot be complex, but by the linearity of the system, if $e^{\lambda t}$ is a solution, then $e^{\lambda t} + e^{\lambda^* t}$ is also a solution, so we can use this as the real solution for $x(t)$.

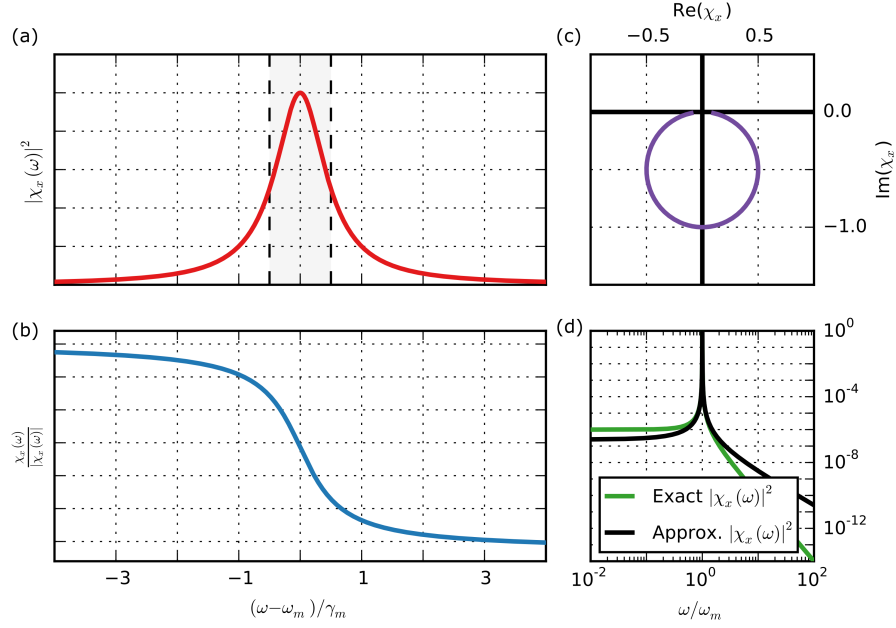


Figure 1.3: Mechanical Resonance: (a) Magnitude (squared) of the complex mechanical transfer function, $\chi_x(\omega)$. (b) Phase of $\chi_x(\omega)$. (c) Parametric plot of $\chi_x(\omega)$ in the complex plane. (d) Comparison of the exact and approximate expressions for $|\chi_x(\omega)|$

This approximated first-order resonant response is equivalent to having described the mechanical motion by a complex spring constant – a model which is sometimes used from the beginning, particularly when describing internal damping. The two models are negligibly different, as long as we only care about the behavior near-resonance³. The exact and approximated response functions are both plotted in Fig. 1.3.

1.3.3 Rewriting as a complex amplitude

For some applications, it is convenient not to work with the real position $x(t)$, but rather to introduce a complex variable c to describe the oscillator. As a first step, we can reduce our second-order equation (Eq. 1.5) to a pair of coupled first-order coupled equations by defining $p = m\dot{x}$, giving:

3. In fact, the relative deviation between the exact and approximate mechanical response functions will be of order $1/Q$ near resonance. For experiments in which the DC or far-off-resonant response of the oscillator is relevant, this approximation is not valid

$$\begin{cases} \dot{x} &= p/m \\ \dot{p} &= -\gamma p - m\omega_m^2 x + F \end{cases} \quad (1.9)$$

Now, we can reduce this further to a single differential equation by moving to a complex amplitude c , defined as

$$c = \frac{1}{\sqrt{2m\hbar\omega_m}}(m\omega_m x + ip) = \frac{1}{2x_{\text{ZPF}}} \left(x + \frac{ip}{m\omega_m} \right) \quad (1.10)$$

where $x_{\text{ZPF}} = \sqrt{\frac{\hbar}{2m\omega_m}}$ is the usual size of the zero-point-fluctuations, as defined in quantum mechanics. This definition is chosen such that $|c(t)|^2 = \frac{E}{\hbar\omega_m}$, where E is the total energy of the oscillator, equal to the sum of kinetic and potential energies. The normalization by $\hbar\omega_m$ is not necessary, but is a standard practice in the field, as it allows for easy translation to a quantum description of the oscillator. Aside from the normalization constant, $c(t)$ simply corresponds to the complex phasor for the mechanical oscillator, which rotates at $+i\omega_m$. The position and momentum of the oscillator correspond (up to a coefficient) to the real/imaginary components of this phasor:

$$x = x_{\text{ZPF}}(c + c^*) \quad (1.11)$$

$$p = -im\omega_m x_{\text{ZPF}}(c - c^*) \quad (1.12)$$

By substitution and algebra, this new variable allows us to transform Eqs. 1.9 into:

$$\begin{cases} \dot{c} = -i\omega_m c - \frac{\gamma}{2}c - \frac{\gamma}{2}c^* - \frac{iF}{\sqrt{2m\hbar\omega_m}} \\ \dot{c}^* = +i\omega_m c^* - \frac{\gamma}{2}c^* - \frac{\gamma}{2}c + \frac{iF}{\sqrt{2m\hbar\omega_m}} \end{cases} \quad (1.13)$$

At this point, we will again make the assumption that $\gamma \ll \omega_m$, or equivalently that the mechanical quality factor, $Q \equiv \omega_m/\gamma \gg 1$. This high- Q approximation allows us to decouple c and c^* , since c will oscillate only in a narrow band around $+\omega_m$ and c^* only around $-\omega_m$. (The c (c^*) terms in the c^* (c) equation represent very off-resonant drives).

If we also redefine the drive term as $c_{\text{in}} = -\frac{iF}{\sqrt{2m\hbar\omega_m}}$, then we can completely describe the behavior of our mechanical mode by the following linear first-order differential equation:

$$\dot{c} = -\left(\frac{\gamma}{2} + i\omega_m\right)c + c_{\text{in}} \quad (1.14)$$

This is the expression which we will use throughout this thesis when we wish to describe our mechanical oscillator. Note that this is precisely the equation we would find via a quantum treatment of the mechanical oscillator, solving for the time dependence of the mechanical annihilation operator \hat{c} and taking the classical limit $\hat{c} \rightarrow \langle \hat{c} \rangle = c$. Indeed, if we wished to describe simple (Gaussian) quantum mechanical effects in the mechanical oscillator, we could use a description like this, while including appropriate quantum noise terms and being careful about commutation and calculation of noise correlators. For our purposes, we will simply work with this classical equation of motion when describing our mechanical motion.

1.4 Optomechanical Coupling

1.4.1 Canonical system

We now consider the generic model of an optical cavity whose frequency is linearly coupled to the position of a mechanical oscillator. This system is quite generic, and describes many experimental realizations, but we'll consider the traditional example depicted in Fig. 1.4: a Fabry-Perot cavity in which one of the mirrors is harmonically bound along the cavity axis. We'll describe the optical cavity and mechanical resonance using the same notation as in Chs. 1.2 and 1.3.

In the absence of any mirror motion, the resonance condition for the cavity is that its length L be divisible by an integer number of optical half-wavelengths: $\lambda/2 = L/n$. Therefore, we can express the bare cavity frequency in terms of its length as $\omega_c = 2\pi c/\lambda = 2\pi nc/L$. Now, if the mirror is displaced by an amount x , such that the new cavity length is $L + x$, the cavity frequency is given by $\tilde{\omega}_c(x) = \omega_c \frac{1}{1+x/L}$, which we can approximate as $\tilde{\omega}_c(x) = \omega_c - Gx$, where $G = \omega_c/L$. Adding this position-dependent frequency to our




Figure 1.4: Canonical Optomechanical System: A harmonically-bound end-mirror (a) results in a position-dependent cavity frequency, shifting the magnitude (b) and phase (c) response of the cavity accordingly. If one considers addressing the unperturbed cavity with an on-resonance laser, its clear that, to first order, the amplitude is insensitive to position fluctuations, while the phase is maximally sensitive. For non-resonant drives, the motion will be transduced to both amplitude and phase fluctuations

equation of motion for the cavity amplitude, we have:

$$\dot{a} = -\left(\frac{\kappa}{2} - i(\omega_c - Gx)\right)a + \sqrt{\kappa_{\text{in}}}a_{\text{in}} \quad (1.15)$$

Or, in terms of the previously defined complex mechanical amplitude:

$$\dot{a} = -\left(\frac{\kappa}{2} - i(\omega_c - Gx_{\text{ZPF}}(c + c^*))\right)a + \sqrt{\kappa_{\text{in}}}a_{\text{in}} \quad (1.16)$$

Note that this approximation of a linear frequency change is typically very good, particularly in macroscopic Fabry-Perot systems like the one considered in this thesis, in which x/L is at most $\approx 10^{-8}$

Above, we found that the coupling constant by which the motion 'pulls' the cavity frequency is $G = \omega_c/L$ for this movable-end-mirror system. Let's now calculate the force on the mirror due to radiation pressure of the cavity light. From electrodynamics, we know that the radiation pressure force of a beam with power P , incident on a perfect reflector is given by $2P/c$. So, to calculate the force due to the intracavity light, we'd like to know the power incident on the mirror. We'll return to our expression for the intracavity field, \bar{a} , and note that $|\bar{a}|^2$ corresponds to the intracavity photon number (i.e. intracavity energy,

normalized by $\hbar\omega_L$). Considering the round-trip time of the cavity ($2L/c$), we see that the power incident on either mirror is

$$P_{\text{cav}} = \hbar\omega_c \times |a|^2 / \left(\frac{2L}{c}\right) \quad (1.17)$$

and thus the force on the mirror is

$$F = \hbar\omega_c |a|^2 / L \quad (1.18)$$

If we insert this force into our complex mechanical equation of motion (Eq. 1.14, using the definition of c_{in}), we find the following:

$$\dot{c} = -(i\omega_m + \frac{\gamma}{2})c - \frac{iF}{\sqrt{2m\hbar\omega_m}} \quad (1.19)$$

$$\dot{c} = -(i\omega_m + \frac{\gamma}{2})c - i\sqrt{\frac{\hbar}{2m\omega_m}} \left(\frac{\omega_c}{L}\right) |a|^2 \quad (1.20)$$

$$\dot{c} = -(i\omega_m + \frac{\gamma}{2})c - ix_{\text{ZPF}} \left(\frac{\omega_c}{L}\right) |a|^2 \quad (1.21)$$

$$(1.22)$$

Noting that $\frac{\omega_c}{L}$ is the coupling G we defined previously, we now have:

$$\dot{c} = -(i\omega_m + \frac{\gamma}{2})c - iGx_{\text{ZPF}} |a|^2 \quad (1.23)$$

At this point, it is standard to define a new coupling rate $g_0 = Gx_{\text{ZPF}}$, which represents the single-quantum coupling rate. This will simplify both the optical and mechanical equations of motion as follows (note that we have included another c_{in} term, to allow for additional forces besides the optical one already considered).

$$\boxed{\begin{cases} \dot{a} &= -\frac{\kappa}{2}a + i\left(\omega_c - g_0(c + c^*)\right)a + \sqrt{\kappa_{\text{in}}}a_{\text{in}} \\ \dot{c} &= -(i\omega_m + \frac{\gamma}{2})c - ig_0|a|^2 + c_{\text{in}} \end{cases}} \quad (1.24)$$

Thus, we have shown that the same coupling constant g_0 governs both the position-

dependence of the cavity frequency and the size of the radiation pressure force. This same correspondence could also be found by writing out a classical (or quantum) Hamiltonian for the system and finding the equations of motion. Moreover, while this derivation of the optical frequency shift and radiation pressure force was specific to the movable-end-mirror system, one finds the same behavior in any dispersively coupled system (in which the mechanical motion changes the effective cavity length).

At this point, we have the coupled equations of motion describing an optomechanical system, which is valid in a variety of implementations. We will next discuss the common method of linearization and solution of this system.

1.4.2 Solving the optomechanical system: Linearization

The system described in Eq. 1.24 is a first-order system of differential equations, with non-linear coupling through the $|a|^2$ term driving \dot{c} . (Had we been working in a quantum picture, this nonlinearity would have been visible in the Hamiltonian, which contains three-wave mixing terms such as $\hat{c}\hat{a}^\dagger\hat{a}$.) However, the size of this nonlinear interaction is set by g_0 , which is generally small compared to other relevant rates in the system. As such, no experiments to-date have observed the nonlinearity of this interaction at the single-photon level. Instead, one typically supposes that the cavity is driven by a large coherent field, which results in an interaction which is now linear, but can be made arbitrarily strong.

We begin by supposing that the cavity is driven by a coherent source with amplitude $\overline{a_{\text{in}}}$ and frequency ω_L :

$$a_{\text{in}}(t) = \overline{a_{\text{in}}} e^{-i\omega_L t} \quad (1.25)$$

This is the same sort of drive we considered in Ch. 1.2, but now to account for the interaction with the mechanical oscillator, we suppose that the cavity field will be given by a mean field plus some small, time-varying fluctuations $d(t)$: $a(t) = e^{-i\omega_L t}(\bar{a} + d(t))$. (Note that $d(t)$ are fluctuations in the frame of the laser.) Ignoring the mechanics, we find the same mean field as before: $\bar{a} = \frac{\sqrt{\kappa_{\text{in}}}}{-\frac{\kappa}{2} + i\Delta} \overline{a_{\text{in}}}$. Plugging this into the full equation, we find that what remains is the following equation for the time-dependence of the field fluctuations $d(t)$:

$$\dot{d} = -\left(\frac{\kappa}{2} - i\Delta\right)d - i\alpha(c + c^*) \quad (1.26)$$

Here, we note the new variable $\alpha = \bar{a}g_0$, which represents the linearized optomechanical coupling rate, enhanced from g_0 by $\bar{a} \propto \sqrt{N_{\text{cav}}}$. Thus, while we have sacrificed the non-linearity of the radiation pressure interaction, the result is a linear coupling rate which can be made arbitrarily large by increasing the strength of the optical drive.

This enhanced coupling also appears in the mechanical equation of motion, which after substituting $a(t) = e^{-i\omega_L t}(\bar{a} + d(t))$ becomes

$$\dot{c} = -\left(\frac{\gamma}{2} + i\omega_m\right)c - ig_0(\bar{a} + d)(\bar{a}^* + d^*) + c_{\text{in}} \quad (1.27)$$

$$\dot{c} = -\left(\frac{\gamma}{2} + i\omega_m\right)c - ig_0(|\bar{a}|^2 + \bar{a}d^* + \bar{a}^*d + d^*d) + c_{\text{in}} \quad (1.28)$$

$$(1.29)$$

The $|\bar{a}|^2$ term is a constant force, resulting in a static displacement which we will ignore. Moreover, since d is assumed to be a small fluctuation relative to \bar{a} , we will ignore the d^*d term. What remains is the following:

$$\dot{c} = -\left(\frac{\gamma}{2} + i\omega_m\right)c - i(\alpha d^* + \alpha^* d) + c_{\text{in}} \quad (1.30)$$

where α is defined as before ($\alpha = g_0\bar{a}$).

We have now reduced our system to a set of coupled, linear, first-order differential equations describing the mechanical amplitude c and the optical fluctuations d :

$$\boxed{\begin{cases} \dot{d} = -\left(\frac{\kappa}{2} - i\Delta\right)d - i\alpha(c + c^*) \\ \dot{c} = -\left(\frac{\gamma}{2} + i\omega_m\right)c - i(\alpha d^* + \alpha^* d) + c_{\text{in}} \end{cases}} \quad (1.31)$$

1.4.3 Solving the optomechanical system: Analysis in Fourier domain

At this point, we can understand much of the physics of this coupled system by moving to the Fourier domain to solve the system defined in Eq. 1.31. Doing so, we have:

$$\begin{cases} -i\omega d[\omega] = -\left(\frac{\kappa}{2} - i\Delta\right) d[\omega] - i\alpha(c[\omega] + c^*[\omega]) \\ -i\omega c[\omega] = -\left(\frac{\gamma}{2} + i\omega_m\right) c[\omega] - i(\alpha d^*[\omega] + \alpha^* d[\omega]) + c_{\text{in}}[\omega] \end{cases} \quad (1.32)$$

Or, rearranging and defining the optical and mechanical susceptibilities:

$$\begin{cases} \chi_c[\omega] = \left[\frac{\kappa}{2} - i(\Delta + \omega)\right]^{-1} \\ \chi_m[\omega] = \left[\frac{\gamma}{2} - i(\omega - \omega_m)\right]^{-1} \end{cases} \quad (1.33)$$

we find

$$d[\omega] = \chi_c[\omega](-i\alpha(c[\omega] + c^*[\omega])) \quad (1.34)$$

$$\chi_m^{-1}[\omega]c[\omega] = -i(\alpha d^*[\omega] + \alpha^* d[\omega]) + c_{\text{in}}[\omega] \quad (1.35)$$

Consider the expression for $d[\omega]$ and note that by replacing $x[\omega] = c[\omega] + c^*[\omega]$, we have:

$$d[\omega] = -i\chi_c[\omega]\alpha x[\omega] \quad (1.36)$$

Thus, we see that the optical fluctuations of the cavity field are proportional to the mechanical motion, scaled by the coupling rate α and filtered by the cavity susceptibility $\chi_c[\omega]$. This illustrates the transduction capability of optomechanical systems: By monitoring the amplitude/phase fluctuations of the light leaving the cavity, one has a direct record of the full spectrum of the mechanical motion.

To better understand the effect of optomechanical coupling on the mechanical motion, we will substitute $d[\omega]$ into the equation for $c[\omega]$, but to do so, we should be careful about $d^*[\omega]$. Specifically, we cannot simply say that $d^*[\omega] = (d[\omega])^*$. Instead, if we were to consider the time-domain equation for d^* and then apply the Fourier transform:

$$\dot{d}^* = -\left(\frac{\kappa}{2} + i\Delta\right) d + i\alpha^*(c^* + c) \quad (1.37)$$

$$-i\omega d^*[\omega] = -\left(\frac{\kappa}{2} + i\Delta\right) d^*[\omega] + i\alpha^*(c^*[\omega] + c[\omega]) \quad (1.38)$$

$$d^*[\omega] = \left[\frac{\kappa}{2} - i(\Delta + \omega)\right]^{-1} (i\alpha^*(c^*[\omega] + c[\omega])) \quad (1.39)$$

$$d^*[\omega] = \chi_c^*[-\omega](i\alpha^*(c^*[\omega] + c[\omega])) \quad (1.40)$$

$$(1.41)$$

With this subtlety accounted for, we can substitute $d[\omega]$ and $d^*[\omega]$ into $c[\omega]$ to find:

$$\chi_m^{-1}[\omega]c[\omega] = |\alpha|^2(\chi_c^*[-\omega](c^*[\omega] + c[\omega])) - |\alpha|^2(\chi_c[\omega](c[\omega] + c^*[\omega])) + c_{\text{in}}[\omega] \quad (1.42)$$

$$\chi_m^{-1}[\omega]c[\omega] = |\alpha|^2(\chi_c^*[-\omega] - \chi_c[\omega])c[\omega] \quad (1.43)$$

$$+ |\alpha|^2(\chi_c^*[-\omega] - \chi_c[\omega])c^*[\omega] + c_{\text{in}}[\omega] \quad (1.44)$$

$$(1.45)$$

Now, we will make the same high-Q argument we made in 1.3.3, wherein we note that $c[\omega]$ will primarily oscillate in a narrow band around $+\omega_m$, and therefore, the c^* terms represent very off-resonant drives, which we can ignore. Therefore, we are left with

$$\chi_m^{-1}[\omega]c[\omega] = |\alpha|^2\left(\chi_c^*[-\omega] - \chi_c[\omega]\right)c[\omega] + c_{\text{in}}[\omega] \quad (1.46)$$

We will define a quantity $\Sigma[\omega]$, which allows us to describe the mechanical motion in terms of a modified susceptibility $\tilde{\chi}_m[\omega]$:

$$\Sigma[\omega] = i|\alpha|^2(\chi_c^*[-\omega] - \chi_c[\omega]) \quad (1.47)$$

$$\tilde{\chi}_m[\omega] = \frac{1}{\frac{\gamma}{2} - i(\omega - \omega_m) + i\Sigma[\omega]} \quad (1.48)$$

$$c[\omega] = \tilde{\chi}_m[\omega]c_{\text{in}}[\omega] \quad (1.49)$$

$$(1.50)$$

$\Sigma[\omega]$ is often called the optomechanical "self-energy", and represents the optical modification of the mechanical resonance. This arises from the self-interaction occurring via the dynamical system formed by the optics/mechanics. From the definition of Σ , we see that it will be a complex quantity depending on the laser power and detuning (through its dependence on α and χ_c). We also note here that the dependence on ω comes only through the cavity susceptibility, which varies with ω on a scale of κ . Since we are considering only a weak interaction, we can assume that the width of the mechanical resonance will be significantly less than κ . Thus, we can assume that $\chi_c[\omega]$ and hence $\Sigma[\omega]$ will be constant with respect to ω over the relevant mechanical bandwidth. Therefore, we can eliminate the frequency dependence and simply evaluate the self-energy at the mechanical frequency: $\Sigma[\omega] \rightarrow \Sigma[\omega_m] \equiv \Sigma$.

By inspecting the appearance of Σ in the effective susceptibility, we see that the real part of Σ corresponds to a shift in the mechanical frequency ("optical spring"), while the imaginary part results in a change in the effective damping rate of the oscillator ("optical damping"). This 'dynamical backaction' is the mechanism by which one is able to control the motion in an optomechanical system, by exploiting and controlling the radiation pressure force. We will now explore this control in more detail, to better understand its limitations and tunability.

1.4.4 Optomechanical Control: Dynamical Backaction via Σ

Since dynamical back-action, set by Σ , forms the basis of optomechanical control, it is worth understanding how it can be tuned by the optical drive, and how this tunability changes based on the system parameters. From the definition $\Sigma = i|\alpha|^2(\chi_c^*[-\omega_m] - \chi_c[\omega_m])$, we recall that $|\alpha|^2 \propto N_{\text{cav}}$, and thus all back-action effects will scale linearly with the intracavity power. Meanwhile, there will be an additional dependence on laser detuning Δ , via the cavity susceptibilities in $(\chi_c^*[-\omega_m] - \chi_c[\omega_m])$. These susceptibilities vary on a frequency scale of κ , and we are evaluating them at the mechanical frequency ω_m , so it is not surprising that the behavior of Σ will be different depending on the relative size of these two parameters. In fact, one typically defines two parameter regimes: the 'resolved-sideband' limit, in which $\omega_m \gg \kappa$, and the 'unresolved-sideband' limit, in which $\omega_m \ll \kappa$.

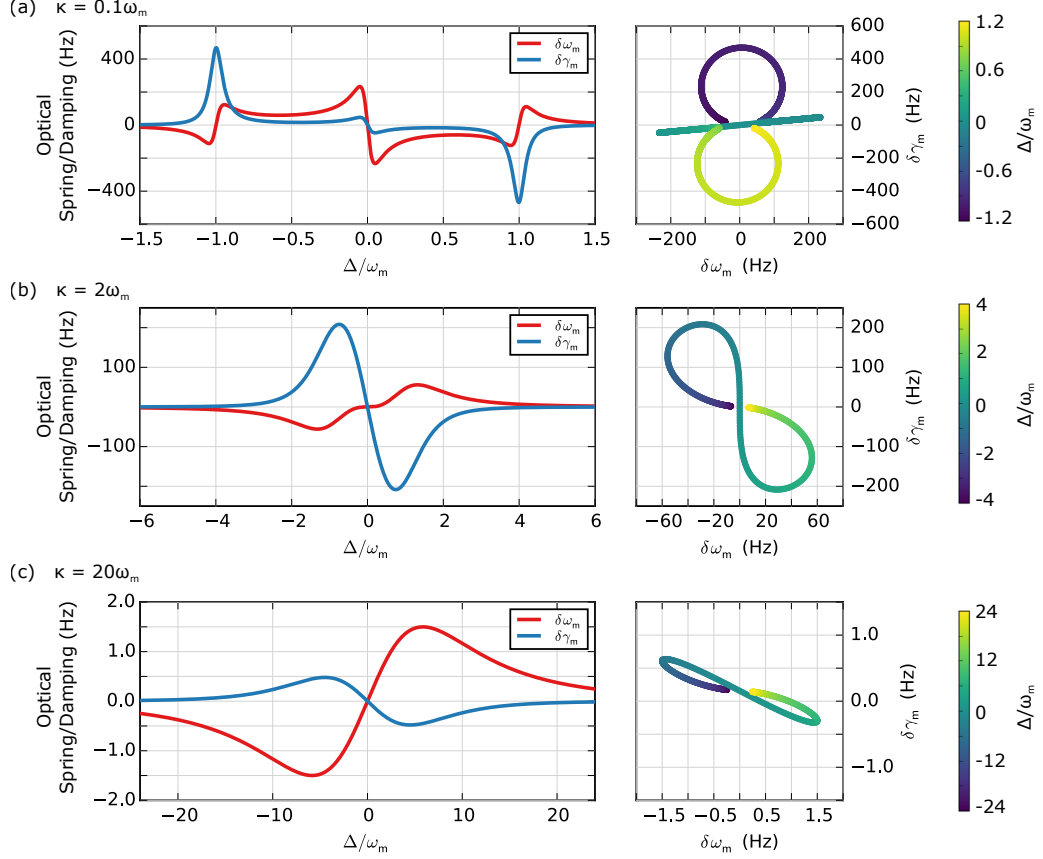


Figure 1.5: Dynamical Backaction in (Un-)Resolved Sideband Regime: Optical spring and damping for different regimes of sideband resolution, plotted separately versus detuning on the left, plotted parametrically in the complex plane on the right. (a) $\kappa = 0.1\omega_m$: Clearly visible interactions for $\Delta = \pm\omega_m, 0$. Positive detunings result in negative $\delta\gamma_m$, which can create a net negative damping rate, leading to self-oscillation. (b) $\kappa = 2\omega_m$: The separated resonance peaks begin to merge. (c) $\kappa = 20\omega_m$: The contributions from the two χ_c terms begin to destructively interfere, pushing Σ onto the real axis.

Optomechanical Control: Resolved-Sideband Limit

When $\omega_m \gg \kappa$, the terms $\chi_c[\omega_m]$ and $\chi_c^*[-\omega_m]$ represent resonant terms occurring at well-separated detunings, $\Delta = \pm\omega_m$. At each of these detunings, only one of the susceptibility terms is significant, and Σ is essentially proportional to that term, which traces out a complex resonance as a function of detuning. We see these resonances at $\pm\omega_m$ in Fig. 1.5(a). There is also a feature around $\Delta = 0$, which comes from the small overlapping tails of $\chi_c[\omega_m]$ and $\chi_c^*[-\omega_m]$. Even though the susceptibility terms are small near $\Delta = 0$, there is still a significant Σ here due to the large intracavity photon number near resonance.

One important characteristic of the resolved-sideband limit is immediately visible from this figure: Detunings near $-\omega_m$ result in damping of the oscillator, while detunings near

$+\omega_m$ result in anti-damping. This resolved-sideband damping is the basis of optomechanical cooling, which has made possible preparation of mechanical oscillators in their ground-state (including, for example, in this group, in [3]).

Since we are considering the full complex behavior of Σ , it can be helpful to view it in the complex plane, plotted parametrically with respect to Δ . This is shown in the right panel of Fig. 1.5(a), where we see the resonant loops at $\pm\omega_m$ as well as the feature near $\Delta = 0$. If one thinks of the mechanical oscillator as having a complex eigenvalue $\omega_m + i\gamma_m$, then the right panels of Fig. 1.5 can be thought of as the modification of the complex mechanical eigenvalue as the laser detuning is varied.

The fact that the features at $\pm\omega_m$ trace out the full χ_c resonance is particularly relevant for optomechanical control. Specifically, if we zoom in on one of these features (say, the feature at $\Delta = -\omega_m$), and recall the overall scaling of Σ with optical drive power, we see that one has nearly-arbitrary control over the mechanical resonance. That is, to achieve any particular $\delta\omega_m$ and $\delta\gamma_m$, one need only choose a particular laser power (to set the size of the loop) and laser detuning (to choose the angle along this loop). This strong control of the mechanical resonance is something which is not possible in the unresolved-sideband regime, as we will see next.

Optomechanical Control: Unresolved-Sideband Limit

When $\omega_m \ll \kappa$, the terms $\chi_c[\omega_m]$ and $\chi_c^*[-\omega_m]$ are resonances of width κ separated by only ω_m , so the features overlap. As such, the exact interference of these terms becomes important. Specifically, one finds that the real parts of $\chi_c[\omega_m]$ and $\chi_c^*[-\omega_m]$ are nearly identical and hence cancel out in Σ . The imaginary parts, however, differ by a sign, and hence do not cancel out in Σ . The net result is that the dynamical back-action described by Σ is primarily real, resulting in a frequency shift, but with no significant optical damping, as seen in Fig. 1.5(c).

The unresolved-sideband behavior of Σ in the complex plane can be seen in the right panel of Fig. 1.5(c). Here, in contrast to the resolved-sideband curve in Fig. 1.5(a), we see that interference of the two resonance loops has pushed the dynamical backaction onto the x-axis. Thus, even if we consider varying the optical power, it is no longer possible to span

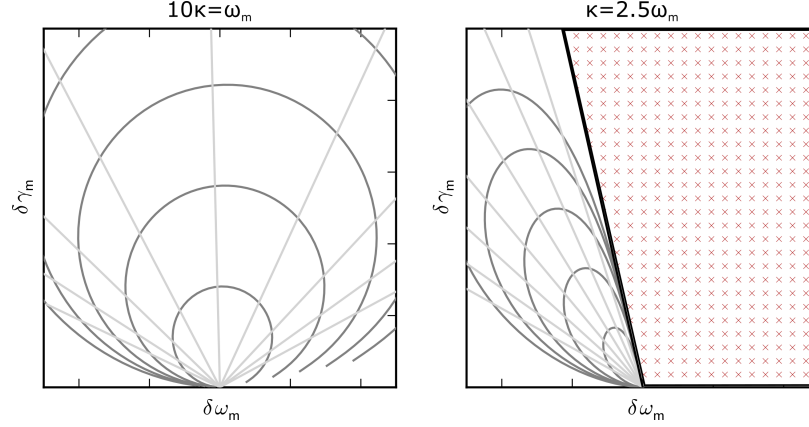


Figure 1.6: Spanning the complex eigenvalue plane. Σ plotted along contours of constant P (light gray) and Δ (dark gray). In (a), $\kappa = .1\omega_m$ (resolved-sideband), and we see that the possible values of Σ span the eigenvalue plane. (We’ve only shown negative detunings, which span the upper half-plane.) In (b), $\kappa = 2.5\omega_m$ (unresolved-sideband), and Σ can no longer access all resonance parameters – the inaccessible region is marked with red hashing.

the $\delta\omega_m, \delta\gamma_m$ plane and access arbitrary values for the mechanical resonance.

This arbitrary control (or lack thereof) in the complex plane is summarized in Fig. 1.6, which shows the complex mechanical eigenvalues along contours of both P and Δ , specifically for negative (red) detunings. Later in this thesis, in which we exploit the tunability of our optomechanical system, it will be important that we are working with a resolved-sideband system.

1.5 Membrane-in-the-Middle Optomechanics

The preceding theoretical description of optomechanical measurement/control applies to any system in which a cavity frequency is linearly coupled to the position of a mechanical oscillator. As discussed before, experimental implementations vary widely in size and geometry. In this work, we focus on the ‘membrane-in-the-middle’ geometry, which involves a thin dielectric membrane dispersively coupled to a Fabry-Perot cavity. Here, we will introduce the theoretical basics of such systems, as well as discuss their role in the optomechanical field as a whole, both historically and looking forward.




Figure 1.7: Membrane-in-the-middle Optomechanics: The dielectric membrane is translated along the cavity standing wave. In regions of minimal (maximum) intensity, the cavity frequency is minimally (maximally) perturbed. Since the dielectric results in a longer effective path length, it can only decrease the cavity frequency from its unperturbed value (dashed gray line)

1.5.1 Model: Dielectric slab within cavity

The basic model for membrane-in-the-middle optomechanics is illustrated in Fig. 1.7. The central idea involves placement of a thin (sub-wavelength) dielectric slab into the standing wave formed in a Fabry-Perot cavity. If the dielectric is placed at the node of the electric field, it will minimally perturb the cavity, but as it is moved towards a cavity anti-node, the index of refraction of the dielectric results in an increased effective cavity length, and hence a lower cavity frequency. As the membrane is moved through the standing wave, one expects an approximately sinusoidal perturbation to the cavity frequency.

To describe the system more carefully, we will make several simplifying assumptions, which are not always sufficient to describe membrane optomechanics, but are completely sufficient for the work presented in this thesis. First, we will consider only a single mode of the cavity, which we model as a plane wave interacting with the dielectric. More detailed models, incorporating higher-order mode profiles and tilted membrane geometry are relevant for some experiments [4, 5], in which the membrane results in a coupling of different

transverse cavity modes, but we need not consider those distinctions here. Instead, we can simply write out the equations for the field amplitudes E_i shown in Fig. 1.8. These different fields are related by mirror reflectivities and dielectric boundary conditions.

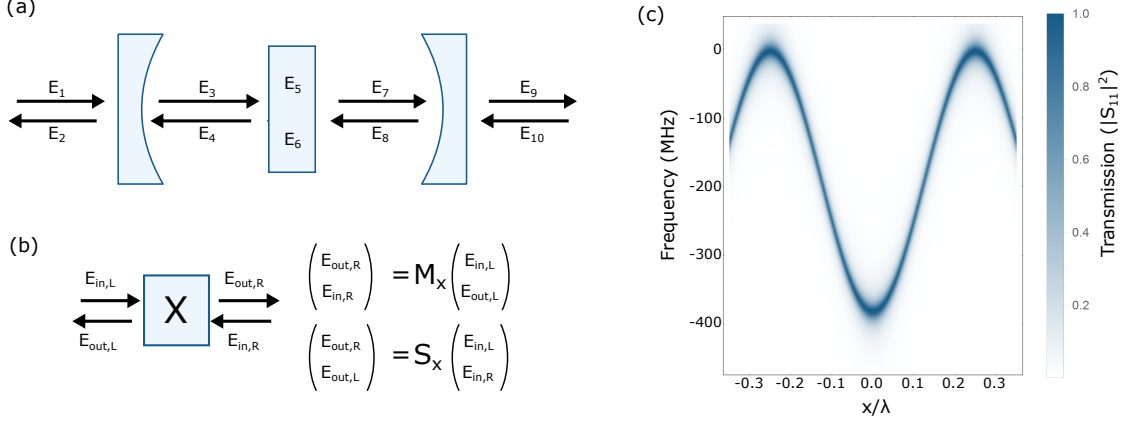


Figure 1.8: Solving for Perturbation of Cavity by a Dielectric Slab: (a) The behavior of the membrane/cavity system can be analyzed by solving the system of equations relating these electric fields, which are linked by boundary conditions. (b) For some purposes, this process can be simplified by using a transfer matrix formalism, which relates input/output fields, similar to a scattering matrix (c) Sample spectrum calculated using this approach.

There are many ways to understand this system, but a convenient, algorithmic solution is offered by the transfer matrix formalism, which seeks to describe the behavior of individual optical elements via simple matrix concatenation. Whereas a scattering matrix \mathbf{S} relates inputs/outputs for an optical element, a transfer matrix \mathbf{M} is reorganized such that it links fields on the left and right, as depicted in Fig. 1.8. In some cases, it is more intuitive to describe an element via its scattering matrix (such as for a mirror, where one has $\mathbf{S} = \begin{pmatrix} t & r \\ r & t \end{pmatrix}$), but one can easily translate between the two approaches:

$$\mathbf{M} = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} = \begin{pmatrix} S_{11}S_{22} - S_{12}S_{21} & S_{12} \\ -S_{21} & 1 \end{pmatrix} \frac{1}{S_{22}} \quad (1.51)$$

By relating fields before/after an optical element, it is straightforward to connect multiple elements in series, by multiplying their transfer matrices. For instance, to describe our membrane-in-the-middle (MIM) system, we can write:

$$M_{\text{MIM}} = M_{\text{mirror}_1} M_{\text{prop:n}_0} M_{\text{n}_0|\text{n}_d} M_{\text{prop:n}_d} M_{\text{n}_d|\text{n}_0} M_{\text{prop:n}_0} M_{\text{mirror}_2} \quad (1.52)$$

where:

- M_{mirror_j} is the transfer matrix for the j -th mirror
- $M_{\text{prop}:n_k}$ describes propagation through a region with index n_k
- $M_{n_1|n_2}$ describes the effect of a dielectric interface between regions of index n_1 and n_2 .

These constituent transfer matrices are given by:

$$M_{\text{mirror}_j} = \begin{pmatrix} \frac{1}{\sqrt{T_j}} & i\sqrt{\frac{R_j}{T_j}} \\ -i\sqrt{\frac{R_j}{T_j}} & \frac{1}{\sqrt{T_j}} \end{pmatrix} \quad (1.53)$$

$$M_{\text{prop}:n_j} = \begin{pmatrix} e^{ikn_j d} & 0 \\ 0 & e^{-ikn_j d} \end{pmatrix} \quad (1.54)$$

$$M_{n_1|n_2} = \begin{pmatrix} \frac{n_2+n_1}{2n_2} & \frac{n_2-n_1}{2n_2} \\ \frac{n_2-n_1}{2n_2} & \frac{n_2+n_1}{2n_2} \end{pmatrix} \quad (1.55)$$

where R_j (T_j) is the j -th mirror power reflection (transmission) coefficient, $k = \frac{2\pi}{\lambda}$ is the wavenumber of the light, d is the propagation distance, and n_0 and n_d are the indices of refraction for vacuum and the membrane dielectric.

Once this full stack has been calculated, it is simple to return to a scattering matrix, suppose that the cavity is driven only from the left side, and compute the reflection and transmission coefficients, S_{21} and S_{11} . A sample transmission spectrum is plotted in Fig. 1.8, as a function of the membrane position and laser frequency. One can also consider the limit of a closed cavity ($R_1, R_2 \rightarrow 0$) in order to analytically calculate the perturbed cavity frequencies [5].

1.5.2 Implementation: Si₃N₄ Membranes

The dielectric slabs used in this thesis (and in most similar research) are high-stress Si₃N₄ membranes, like the one shown in Fig. 1.11. In our lab, these are purchased commercially,

with typical membrane lengths of 0.5-1mm and thickness $\sim 50\text{nm}$. Without any additional modification, these membranes have mechanical quality factors of $\sim 10^6$ at room temperature and $\sim 10^7$ at cryogenic temperatures. Moreover, they have been found to have very low optical absorption (imaginary part of index of refraction $\sim 10^{-4}$ [6]). This means that they can be incorporated into a high-finesse optical cavity without significantly degrading the cavity linewidth. Furthermore, low optical absorption is important to avoid heating of the membrane – in our experiments, we frequently reach $\sim 10\text{mW}$ of optical power passing through the membrane, while still maintaining thermalization to the cryostat.

The vibrational modes supported by a stressed square membrane are depicted in Fig. 1.9, labeled by the number of anti-nodes in each of the $\{x,y\}$ dimensions. The frequency of the $\{i,j\}$ mode is proportional to $\sqrt{i^2 + j^2}$, so relative to the fundamental ($\{1,1\}$) mode, one can expect the frequencies to follow

$$\nu_{i,j} = \nu_{1,1} \sqrt{\frac{i^2 + j^2}{2}} \quad (1.56)$$

If the membrane sides are not equal (say, L_1 and L_2), then the Eq. 1.56 becomes

$$\nu_{i,j} = \nu_{1,1} \sqrt{\frac{(\frac{i}{L_1})^2 + (\frac{j}{L_2})^2}{2}} \quad (1.57)$$

This mode spectrum has an implication which will be important to our later work: For a perfectly square membrane, the $\{i,j\}$ and $\{j,i\}$ modes would be degenerate. Any small deviations from a square membrane will break this degeneracy, but in general we can expect these modes to be nearly-degenerate given standard fabrication precision. Fig. 1.10 shows a sample of the mode spectrum measured for our membrane. The expectations for a square membrane are marked, and we see that the mode spectrum matches the expectation from Eq. 1.56, including the nearly-degenerate mode pairs. We will return to one such mode pair in Ch. 4.

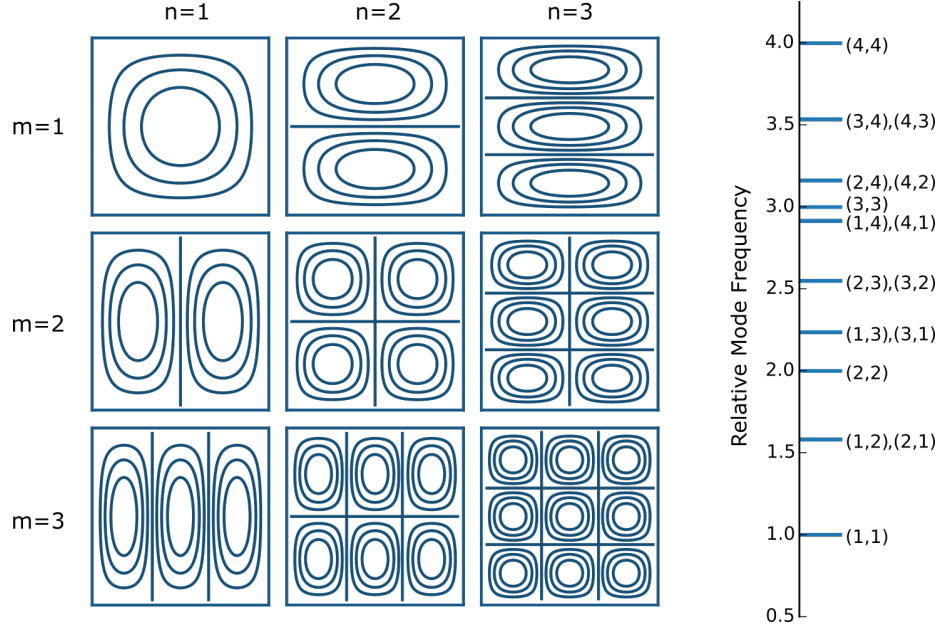


Figure 1.9: Square membrane vibrational mode patterns (a) and frequencies (b)




Figure 1.10: Measured thermal motion spectrum. Red markers indicate prediction of Eq. 1.56, given the fundamental frequency $\nu_{1,1} \approx 352\text{kHz}$. The other, smaller peaks, are generally assumed to be other vibrational modes, either of the membrane support structure or of the cavity mirrors.

1.5.3 Historical context: Membranes at Yale

The use of dispersively-coupled dielectric membranes as optomechanical systems began in this group in 2007. The membranes used were commercial products, available from a MEMs technology company which produced them for use as x-ray transmission windows and TEM sample holders. The membrane-in-the-middle geometry marked a departure from most other work in the community, which was focused on the movable-end-mirror geometry, using reflective coatings fabricated on cantilevers or other vibrating structures. After demonstrating the operation of the system [7, 8] and characterizing the mechanical properties of Si_3N_4 membranes [6], the group studied optical avoided crossings in the system

[9], which were promising for quadratic coupling applications. The next several years were spent developing a cryogenic version of the experiment, with the goal of cooling one of the membrane’s vibrational modes near its quantum ground state, using a combination of passive cryogenic cooling and active laser cooling. This system was limited by a number of technical issues, including environmental acoustic noise and classical laser noise, but was able to demonstrate cooling to a mechanical occupancy of ~ 30 in 2011 [10]. The work in this thesis began in early 2012, during the transition to a 2nd-generation cryogenic system, redesigned to resolve various technical problems, and make possible experiments like ground-state cooling.

Early work in my PhD involved construction of a new cryogenic platform for the optomechanical device, and modification of the optical setup to reduce laser noise. The cryogenic platform was designed around the idea of mechanical stability, as well as tunability of the optical/mechanical alignment. Rather than sending light into the cryostat via free-space, this new generation relied on input via an optical fiber, which introduced the need for cryogenic motors for aligning the beam to the cavity mode. In addition to various mechanical redesigns within the cryostat, the optical setup was also modified to include a second filter cavity for reducing noise on the measurement laser. The introduction of this cavity would also necessitate some redesign of the various feedback systems used in the setup. This experimental setup is described in Ch. 2, and additional details are available in [5].

1.5.4 State-of-the-field: Membranes elsewhere

Since their introduction (as optomechanical devices) at Yale in 2007, Si_3N_4 membranes have been adopted by numerous groups, with applications spanning the breadth of optomechanics. While summarizing all of the ongoing work in membrane optomechanics is infeasible, what follows is a brief overview of some of the major progress, to help identify the context of this experiment.

One of the early advantages of Si_3N_4 membranes was the high mechanical quality factors ($Q \approx 10^7$) available in a cheap, commercially-available product. Some of the progress in membrane optomechanics has been focused around improving these Q factors, as well as making them less sensitive to mounting conditions. Particular success has been found via




Figure 1.11: Standard and engineered Si_3N_4 membranes. (a) A standard Norcada membrane, held by tweezers. (b) A membrane with a phononic shield patterned in the Si chip (from Regal group: [11]). (c) A membrane in which the Si_3N_4 itself has been patterned to create a localized vibrational mode with low curvature at its boundaries (from Schliesser group: [12])

phononic engineering, both of the Si substrate which supports the membrane, as well as of the membrane itself. Early work [11] from the Regal group demonstrated that by carving a phononic ‘shield’ in the Si chip, one could engineer an acoustic bandgap that suppressed non-membrane mechanical modes, eliminating unwanted mechanical backgrounds and improving membrane mode Q -factors.

Recent work has also dramatically improved the mechanical Q -factors available with Si_3N_4 membranes, by defining the mechanical mode in a way that reduces intrinsic dissipation through minimization of the mode curvature [12]. This ‘soft clamping’ of the mechanical mode has achieved impressive mechanical Q -factors, approaching 10^9 , opening up a new parameter regime in which the optomechanical interaction rate can be made significantly larger than the thermal decoherence rate.

In addition to pushing membrane optomechanical systems towards higher Q -factors, many in the community have sought to develop hybrid optomechanical system by simultaneously coupling in additional degrees of freedom. A joint effort by the Regal and Lehnert groups [13] has sought to achieve simultaneous optomechanical coupling of a membrane to an optical cavity as well as an LC microwave resonator. Such efforts aim to exploit the mechanical element as an intermediary, to transfer (quantum) information between the microwave and optical domains – a critical goal for optically linking superconducting microwave qubit systems. Other ‘hybrid’ systems involve coupling of membranes to spins, magnetic fields, and atomic ensembles.

1.5.5 Previous Optomechanics Work with this System

The focus of this thesis will be multimode physics and adiabatic dynamics associated with exceptional points, but this in fact constitutes only the latter half of my PhD work. Prior to studying interactions of multiple mechanical modes, we carried out two major experiments using the system described in Ch. 2. In the first, we studied an extension of traditional optomechanics, in which two optical modes are coupled via a common interaction with the membrane. Specifically, we worked in a regime in which two cavity modes undergo an avoided crossing, as the membrane position is varied (see Fig. 1.12). If the membrane is statically placed exactly at this avoided crossing, the cavity frequencies will have a quadratic dependence on the membrane position, which leads to a qualitatively different form of optomechanical interaction. Our primary result was the demonstration of the dynamical backaction (optical spring/damping) arising from this quadratic coupling. The availability of multimode interactions and tunable linear/quadratic coupling remains a key advantage of this experimental system, compared with similar setups in other groups.




Figure 1.12: Optical avoided crossing from [14], depicting quadratic optomechanical coupling.

The next major thrust of my PhD work was an attempt to cryogenically- and laser-cool one of our vibrational modes near its quantum ground state. This represented the major outstanding goal from the previous generation, and required resolution of a number of technical challenges. Foremost among these was laser noise, which sets a fundamental limit on laser cooling, by introducing an additional source of force noise on the membrane. We overcame this by use of separate filter cavities for each laser, filtering noise at ω_m by

$\sim 30\text{dB}$. Keeping these cavities locked simultaneously with the experimental cavity presented another challenge, which was overcome by implementing a feedback scheme in which most of the feedback signal is applied to a modulator placed after the filter cavities (see Ch. 2). This scheme continues to be a particularly convenient feedback scheme, even when the filter cavities are not in use. By use of filter cavities and low-noise RF circuitry, we were able to produce a cooling beam which was dominated by quantum noise, even at the large powers relevant for ground-state cooling. Using a heterodyne measurement scheme, we were able to separately record both the red and blue mechanical sidebands, for a range of cooling powers. This simultaneous measurement of both sidebands allowed us to very cleanly look for a quantum asymmetry which is only visible at low phonon occupancies. We analyzed the extracted thermal sidebands to infer the effective temperature, using multiple semi-independent analytic techniques as cross-checks. In the end, our coldest data was found to correspond to an effective mechanical occupation number of $n_m = 0.84 \pm 0.22$, with consistency between all 4 analysis techniques, as summarized in Fig. 1.13. The achievement of this low occupation number was competitive with some of the coldest oscillators to-date at the time, particularly due to its large mass (and low frequency).

After the completion of these projects, our attention turned to optically-mediated interactions between mechanical modes, which eventually led us to the study of degeneracies in damped systems and exceptional points, which will be the focus of this thesis.

I will begin by introducing our experimental system in Chapter 2, then provide a pedagogical introduction to exceptional points, via a simplified theoretical model, in Chapter 3. Our first main experimental results, demonstrating key features of exceptional points in an optomechanical system, will be presented in Chapter 4. Finally, Chapter 5 will describe, both theoretically and experimentally, a subsequent protocol which allows us to expand on the results from Chapter 4.




Figure 1.13: Resolved-sideband laser cooling of a 705kHz membrane mode, from [3]. Low phonon occupancy (d) is confirmed by linewidth thermometry (b), calibrated integration of the red/blue sidebands (c), and sideband asymmetry (d).

Chapter 2

Experimental Setup

We'll now review the experimental setup used for the work presented in this thesis. This includes the cryogenic device (the cavity, membrane, and associated support hardware), as well as the optical and electrical systems used to interrogate the optomechanical system. Additional details regarding specific experimental protocols will be presented in the relevant chapters.

2.1 Cryogenic Optomechanical Platform

The cryogenic device used here was constructed during the first year of my PhD, based on a design intended to solve various technical problems from a previous generation of the experiment. The leading design consideration was vibrational stability of the cavity/membrane system, as prior work had suffered from vibrations originating both inside and outside the cryostat. These environmental vibrations can add unwanted noise to the measurement of the mechanical motion, as well as introduce cavity frequency noise which can be difficult for the laser feedback to track. Additional design considerations included in-situ alignment of the optical beam to the cavity mode, tunability of the membrane orientation, and strong thermal anchoring of the membrane to the cryostat.




Figure 2.1: Schematic of the major components of the ^3He cryostat.

2.1.1 Cryostat

To reduce the thermal occupation of the membrane, the system is housed in a ‘wet’ ^3He cryostat (schematic in Fig. 2.1). The operation of this cryostat relies on a multi-stage system involving evaporative cooling of separate reservoirs of ^3He and ^4He . The outermost layer of the cryostat is a vacuum-jacketed ^4He bath, into which is submersed an inner vacuum chamber (IVC) containing the refrigeration and experimental device. The helium bath space is vented to the atmosphere and replenished approximately weekly by transferring from an external source. This replenishable bath of ^4He distinguishes this system from ‘dry’ cryostats, which operate exclusively using a closed cycle of continuously-recondensed ^4He (thus requiring continuous use of pumps/compressors which cause significant vibrations). Within this IVC is a chamber of liquid ^4He , which is fed by the outer ^4He bath and can either be vented to atmosphere or connected to an external pump. Pumping this chamber (the ‘1K pot’) reduces its temperature to $\sim 1.2\text{ K}$ via evaporative cooling. This 1K reservoir can now be used to cool another isolated chamber of gaseous ^3He located within the IVC. By pumping/cooling the 1K pot, one triggers condensation of the gaseous He-3, forming a volume of liquid ^3He in the coldest section of the cryostat. A high-surface-area material

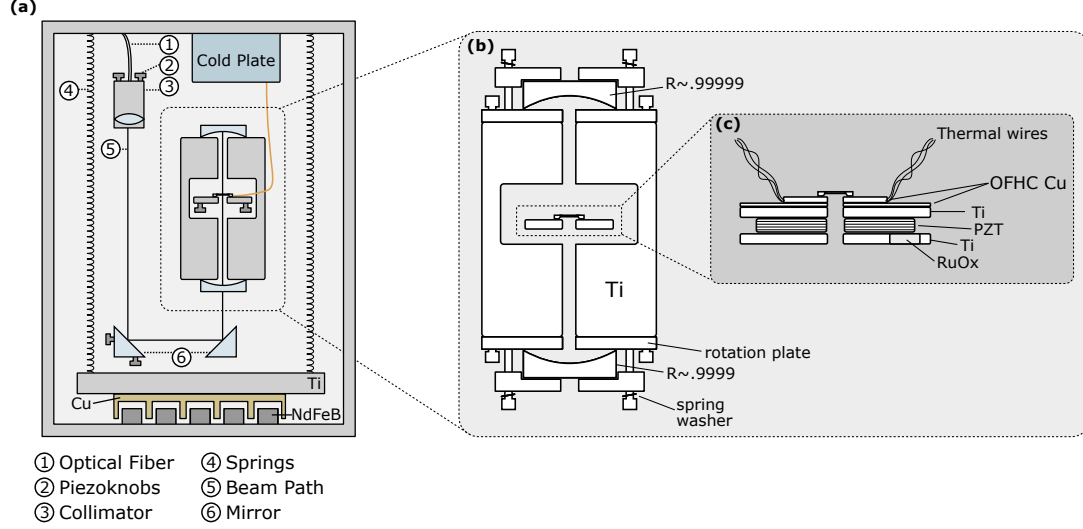


Figure 2.2: Cryogenic Membrane-in-the-Middle Setup. (a) Overview of cryogenic platform. The collimator, angled mirrors, cavity, and membrane are all mounted on the suspended Titanium platform. Knobs on the collimator, angled mirror, and membrane represent piezoelectrically-controlled mirror mounts, capable of tip/tilt adjustment, as well as translation. Thin thermal wires link the cold plate of the cryostat to the membrane support, as well as other parts of the platform. Copper fins on the platform nest between Neodymium magnets (not attached to the platform) for eddy-current damping

in the ^3He space (charcoal) results in adsorptive pumping of the ^3He , further reducing its temperature to a ‘base’ temperature of ~ 300 mK. This evaporatively-cooled ^3He liquid will persist for ~ 1 week in our system, based on the heat capacity of ^3He and the incident heat load of the experimental apparatus. Once the liquid has entirely evaporated and is trapped in the charcoal ‘sorb’, the cryostat returns to 4 K. At this point, one can resistively heat the charcoal ‘sorb’ to ~ 45 K, causing the ^3He to become gaseous again, allowing for recondensation via pumping the 1K pot.

The entire cryostat is suspended on a pneumatically-floated frame, to provide seismic isolation. This required design/installation of a custom support frame to suspend the cryostat from above, rather than sitting the dewar on the floor, as had been done in the past.

2.1.2 Optical Cavity

The optical cavity is formed by two high-reflectivity mirrors mounted on a 3.7cm-long Titanium spacer, as shown in Fig. 2.2(b). The input and output mirrors have reflectivities of 0.9999 and 0.99999, respectively, to create a (nearly) ‘single-sided’ cavity, in which light leaking out of the cavity predominantly returns along the input path. The mirrors are

formed by a stack of alternating dielectric coatings, deposited on a glass mirror substrate by ATFilms. The mirrors sit directly on the cavity spacer, and are held in place by a clamping plate and screws with spring washers. The spring washers ensure that even in the presence of nonuniform thermal contraction, the mirror is still held in place, and without too much force. A single layer of Kapton[®] tape separates the mirrors from the spacer and clamp plates, in order to cushion the interface between the glass and metal. The 3.7 cm cavity length corresponds to a free-spectral range of $\sim 4\text{GHz}$, and the measured cavity decay rate is $\kappa = 180\text{kHz}$ for the work presented here. This decay rate can either be found by measuring the cavity linewidth (via the reflection dip) or by measuring optomechanical effects, such as optical spring and damping.

2.1.3 Membrane

The membrane is a $1\text{mm} \times 1\text{mm} \times 50\text{nm}$, high-stress Si_3N_4 membrane, manufactured by Norcada. The only post-purchase processing applied to the membrane is cleaning via a sequence of acetone/methanol/isopropanol/ O^2 -plasma. The membrane used in this work was selected from a batch of similar membranes based on its room-temperature quality factor, which was on the order of 10^6 . At base temperature, we find membrane quality factors up to 20×10^6 .

These Q values are measured via mechanical ringdowns – that is, we drive (optically or via piezo) the mode to some large amplitude, then record its motion as it decays. Since the energy of the oscillator decays as $e^{-\gamma t}$, the amplitude of motion should decay as $e^{-\gamma t/2}$. Fig. 2.3 shows several example ringdowns, for the first 3 symmetric modes of the membrane.

The membrane is mounted on a multi-layer support (Fig. 2.2(c)) for insertion into the optical cavity. The membrane’s Si chip is held on a circular plate made of oxygen-free high-conductivity Copper, using small amounts of epoxy¹ at three of its corners. Another thin sheet of copper lies underneath this circular plate, for clamping of thermal anchoring wires (described in the next section). The membrane and copper layers are screwed onto a Titanium plate, which is mounted on top of a ring-shaped piezo actuator, which provides

1. Stycast[®] 2850 epoxy. This epoxy is frequently used in the cryostat, due to its high thermal conductivity, low thermal expansion, and compatibility with cryogenic temperatures




Figure 2.3: Examples of mechanical ringdowns for the $\{1,1\}$, $\{2,2\}$, and $\{3,3\}$ modes at frequencies $\omega_{\{1,1\}}=352$ kHz, $\omega_{\{2,2\}}=705$ kHz, and $\omega_{\{3,3\}}=1057$ kHz (plotted in red, blue, and green, respectively). The dashed black lines are fits of the decaying amplitude to $\sqrt{(ae^{-\gamma t/2})^2 + b^2}$ (where a is the initial amplitude of the motion and b is a background). Q is calculated via $Q = \omega/\gamma$

translation of the membrane along the cavity axis. Finally, this piezo sits on another titanium plate, which mounts to a piezo-electrically-controlled mount (described in the next section).

At cryogenic temperatures, we find $\sim 200\text{nm}$ of translation range from the ring piezo. During most experiments, we simply want to keep the membrane at a fixed position, and thus apply a DC voltage. To avoid injecting any force noise onto the membrane, it is important that we minimize voltage noise on this piezo. To achieve this, we supply the voltage from batteries², as well as use aggressive low-pass filtering³

2.1.4 Inside the Cryostat

The cavity, membrane, and supporting hardware are situated on a custom-built Titanium platform inside the IVC, illustrated in Fig. 2.2(a). Titanium was chosen for its low thermal contraction and high thermal conductivity at cryogenic temperatures, and was used

2. To reach the full extension range of the piezo requires $\sim 200\text{V}$, which we create via series connection of $\sim 20 \times 9\text{V}$ batteries.

3. The membrane piezo is essentially 50 nF capacitor to ground, so adding a $1\text{ M}\Omega$ resistor in series before the piezo results in a $\sim 1\text{ Hz}$ low-pass filtering of the voltage across the piezo.




Figure 2.4: Photos of experimental setup. (a) Cryogenic platform inside IVC. (b) Membrane support. (c) Optical cavity, with membrane in place.

throughout the apparatus to ensure uniform contraction of the various components. The platform is suspended on springs, to provide isolation from environmental vibrations (from both inside and outside the cryostat). Eddy-current damping is used to critically damp the springs, via copper fins on the platform and permanent magnets mounted underneath.

Light enters the cryostat via a single-mode fiber⁴, so in order to couple light exiting the fiber to the cavity mode, it is necessary to have collimation and alignment optics inside the cryostat. The optical path therefore includes a collimator and steering mirror, which are both mounted on 3-axis piezo-electrically-controlled mounts. This provides the tunability necessary for aligning the beam to the cavity mode. The piezoelectrically-controlled mounts are manufactured by Janssen Precision Engineering, and are structurally similar to traditional 3-axis optical mounts, in which the orientation of the optic can be controlled by adjusting three screws which tilt/tip the mount. The motion of all three screws also provides translation. In the Janssen design, these screws are turned via a piezo-electric stick-slip system. These motorized mounts are vacuum- and cryogenic-compatible, and we have found them to be stable with respect to long-term position drift. Janssen also consulted on the design of the cryogenic platform which holds these various motors and cavity components.

The titanium ‘bridge’ holding the membrane is also mounted on one of these piezo-electrically-controlled mounts, to allow for long-distance ($\sim 2\text{mm}$) translation of the membrane along the cavity axis (since the ring piezo allows for only $\sim 200\text{nm}$ of translation) as well as control the orientation (tip/tilt) of the membrane relative to the cavity axis. This control was particularly important for projects related to optical avoided crossings [14].

Other features of the cryogenic platform include two RuO_2 thermometers (one attached to the Titanium base, the other to the membrane ‘bridge’), and electrical contact sensors for ensuring that the suspended Titanium platform is not in contact with the walls of the IVC. There is also a photodiode mounted above the cavity output mirror, though this was generally not used due to its low responsivity at cryogenic temperatures.

4. This is in contrast to the previous generation, which sent light in via free-space, travelling $\sim 1.5\text{m}$ through a narrow tube before reaching the cavity. This free-space system suffered from alignment difficulties, particularly due to vibrations of the cryostat.

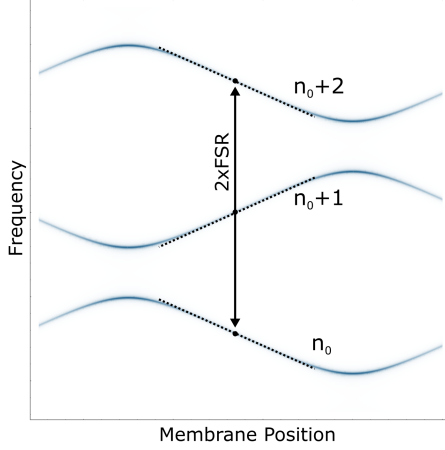


Figure 2.5: Position-dependence of different longitudinal modes, for a membrane located near the cavity waist. The important feature here is the opposite slope of alternate longitudinal modes.

It is important to maintain a strong thermal link between the cold plate of the cryostat and the various parts of the experimental apparatus, particularly the metal plate supporting the membrane. At the same time, it's also important to maintain good vibrational isolation between the cryostat and the membrane, so this thermal connection should not be rigid. To meet these goals, we use a large number ($\sim 10^3$) of very thin, gold-coated copper wires, grouped loosely, as seen in Fig. 2.4.

2.2 Optical Setup

The general optical scheme used in most of the experiments carried out here is depicted in Fig. 2.6. We work with two different cavity modes, separated in frequency by two FSR. As illustrated in Fig. 2.5, cavity modes separated by $2 \times \text{FSR}$ will experience the same cavity shift⁵ with respect to membrane position (assuming the membrane is located near the center of the cavity). By utilizing two cavity modes, we can use separate, non-interacting optical drives for measurement and control.

The measurement beam consists of 4 tones, as seen in Fig. 2.6(a). The main measurement beam is locked on-resonance via a Pound-Drever-Hall scheme (using phase-modulation sidebands at $\pm 15\text{MHz}$ – discussed in Ch. 2.3). A much larger local oscillator (LO) is located

5. Technically, the slope $\frac{\partial \omega_c}{\partial x}$ is slightly different for these two modes, but for a particular position halfway between the node and anti-node of the cavity field, the slopes will be equal. We always operate at this point.

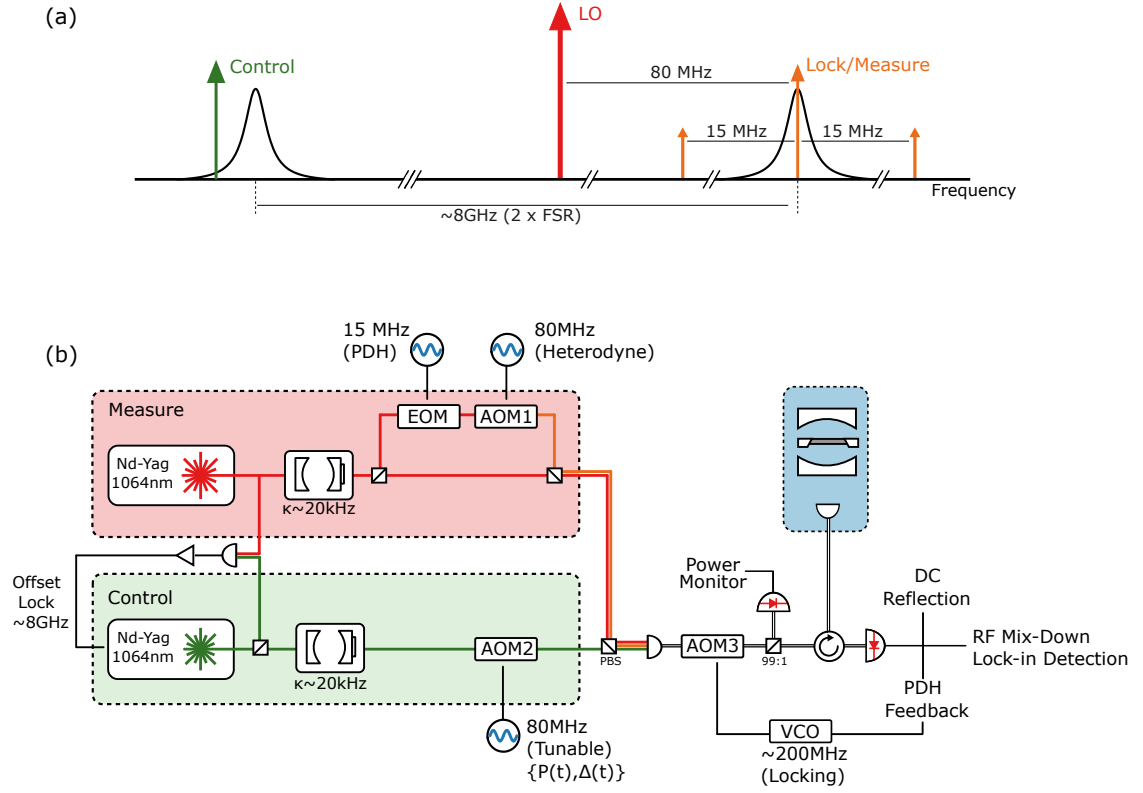


Figure 2.6: (a) Optical beam arrangement in frequency space. (b) Optical circuit schematic.

80MHz away, to allow for heterodyne detection of the sidebands acquired by the measurement beam. The measurement beam is typically between $\sim 5\mu\text{W}$ and $40\mu\text{W}$, depending on whether the goal is to minimize its backaction or maximize the measurement signal-to-noise. Its detuning is typically kept around -10kHz (a small fraction of κ), to minimize its backaction on the mechanical oscillator. The measurement beam must be kept slightly red-detuned, since the high Q of our mechanical modes implies that 1-2 Hz of negative (anti-)damping will result in a negative-linewidth oscillator, leading to self-oscillations and large (un-trackable) cavity fluctuations.

The control beam is (typically) a single beam, locked a fixed frequency ($2 \times \text{FSR} \approx 8\text{GHz}$) away from the measurement beam. By controlling the fixed offset frequency, its detuning relative to its cavity mode can be adjusted (again, generally constrained to negative detunings). The precise details of how the measurement and cooling beams stay locked will be discussed in the feedback section (Ch. 2.4).

To create these beams, we start with two physically distinct lasers (Prometheus Nd-YAG

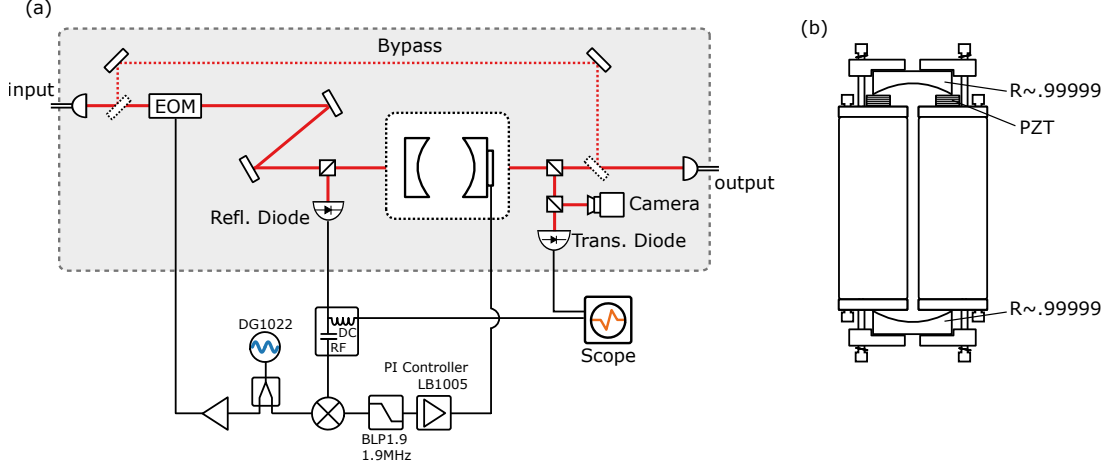


Figure 2.7: (a) Optical/electrical schematic of a filter cavity. (b) Design of the filter cavity, which is similar to the experimental cavity, but without the membrane and with the addition of a piezo (PZT) for adjusting the cavity length.

lasers by Innolight⁶) with $\lambda \approx 1064$ nm. These are low-noise lasers (1kHz linewidth), which include a built-in active intensity stabilization (‘noise eater’) feature. Their frequencies can be shifted either thermally (slow, coarse control over ~ 60 GHz) or using the laser piezo (fast, kHz-scale control over a range of 400MHz). The output power of each laser is ~ 1 W, which allows for losses throughout the optical circuit while still providing sufficient power for interacting with our cavity.

Immediately after the laser outputs, a small fraction of each beam is selected using a beamsplitter, and the two samples are combined on a high-speed photodiode. This beat note near 8GHz will be used to stabilize the relative frequency of the two lasers⁷.

Both lasers then pass through room-temperature, narrow-linewidth ($\kappa \sim 20$ kHz) Fabry-Perot cavities, which act as filters (in the same way that LC resonance circuits can be used as bandpass filters). These cavities are based on the same design, which includes a piezo actuator between the mirror and cavity body, allowing for kHz-scale feedback to the cavity length/frequency. A schematic of the filter cavity is seen in Fig. 2.7(a), including components for alignment, feedback, and an alternate optical path which allows us to bypass the filter cavity when desired.

6. Innolight is now Coherent

7. Details in Ch. 2.4

After passing through its filter cavity, the measurement beam passes through a beam-splitter, with one arm receiving additional modulation via an electro-optic modulator (EOM) and an acousto-optic modulator (AOM). The EOM applies phase modulation at $\sim 15\text{MHz}$ for the PDH lock, and the AOM shifts these beams 80MHz away from the unaltered LO path. The control beam also passes through an AOM, which allows us to control its power and frequency (over a range of several MHz). For some of our experiments, it will be relevant that we can dynamically tune the control beam power/detuning, by modulating the RF source driving this control AOM.

The measurement and control beams are combined into a single fiber via a fiber-based polarizing beamsplitter⁸. Both beams then pass through a fiber-coupled AOM, which shifts the beams by $\sim 200\text{MHz}$. This AOM is used to shift all the laser beams equally as part of the cavity feedback scheme (described in Ch. 2.4). Choosing a fiber-coupled AOM with a high frequency allows us to maximize the bandwidth over which we can shift the laser. (When using AOMs as a frequency shifter, there is an accompanying proportional spatial shift. When this light is then coupled into a single-mode fiber, this spatial shift translates into an attenuation, which determines the usable bandwidth of the AOM. In a fiber AOM, the fiber-coupling happens very close to the AO crystal, which minimizes this effect. Also, since the usable bandwidth will be some fraction of the overall shift, choosing a high-frequency AOM is helpful).

After the AOM, $\sim 1\%$ of the beam power is redirected via a beamsplitter (for monitoring the beam powers), then all beams are directed to the cryostat via a fiber circulator. Finally, this circulator redirects all reflected light to a low-noise photodiode for heterodyne detection.

2.3 Detection Electronics

The reflected light from the cavity is incident on an InGaAs photodiode (PDA10CF). This is an amplified photodiode, with low dark noise and a 150MHz bandwidth (the highest-frequency relevant signal will be the heterodyne beat note near 80MHz). Since photodiodes

8. The relative polarization of the control/measurement beams is not important, but using a PBS allows us to efficiently combine them with minimal loss

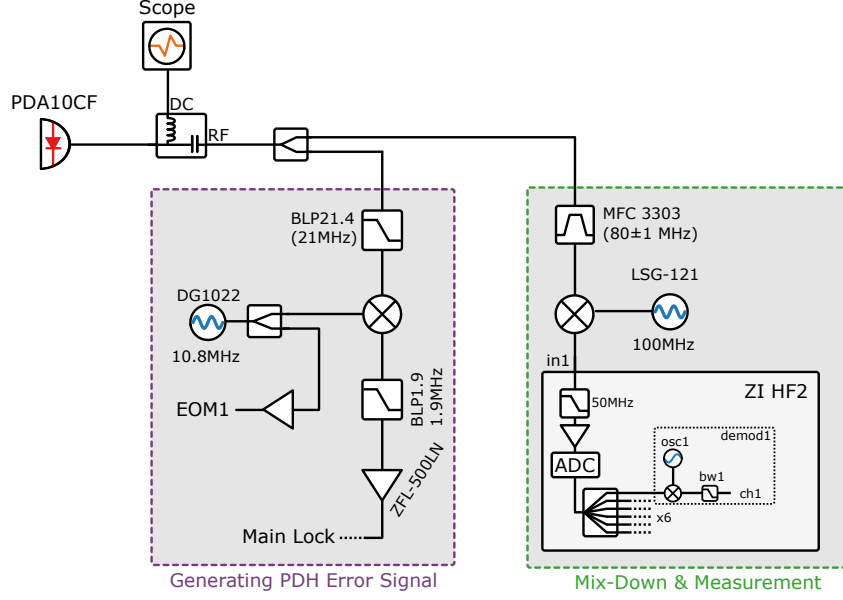


Figure 2.8: Detection and RF processing of photocurrent.

are power detectors, the photocurrent will include a DC component proportional to the total power, as well as AC components coming from beat notes between all pairs of optical tones. Given the optical tones in Fig. 2.6(a), the photocurrent will include many components, with the most relevant being two 15MHz beat notes (which create the PDH error signal) and the 80MHz heterodyne beat note. Any sidebands on the measurement beam will appear as sidebands on the heterodyne beat note.

The electrical circuit depicted in Fig. 2.8 serves to separate the various components of the photocurrent. First, a bias tee separates the DC voltage from any RF signals ($>10\text{MHz}$), then the RF path is split for feedback and measurement. The feedback path is mixed with a 15MHz reference signal (with the mixed-up 30MHz signal removed by a filter), then amplified to form the PDH error signal. The 80MHz component of the measurement path is isolated using a bandpass filter, then mixed-down via a 100MHz reference signal, forming a 20MHz signal suitable for measurement via a lock-in amplifier (Zurich Instruments HF2).

2.4 Feedback Techniques

Keeping all of the lasers and resonators frequency-locked is a major technical challenge in this work, and the experiments described here often involve simultaneous use of 5 inter-dependent feedback circuits. For most (but not all) of these locks, the goal is to stabilize the detuning of a laser relative to a cavity. In principle, this can be done by a variety of methods, all relying on the generation of an ‘error signal’ which varies linearly with cavity detuning. We choose to employ the Pound-Drever-Hall (PDH) technique (a common and versatile form of cavity feedback).

PDH locking is based on the idea that the cavity phase response provides a good linear error signal for detunings near zero. In order to see the phase response of a beam interacting with a cavity, it is necessary to reference it against some other beam which is not interacting with the cavity. (Indeed, this is the basis of using a local oscillator in heterodyne/homodyne detection of phase information.) In principal, one could use the phase response of the heterodyne beat note, but extracting this phase to produce a usable error signal involves non-trivial digital/analog processing which can be unstable with respect to optical phase drifts. The PDH technique essentially uses the interference of two such beat notes to generate the error signal directly, in a way that is only sensitive to cavity frequency changes, and not other phase drifts.

To generate a PDH lock, one begins by phase-modulating the locking laser at a frequency f_{mod} , which is usually much greater than the cavity linewidth. These phase modulation (PM) sidebands will each beat with the carrier beam on the photodiode, generating two signals at f_{mod} which (usually) destructively interfere. As either the carrier or the sidebands is detuned near the cavity, its reflected phase changes, altering the phase of the beat notes and generating a non-zero f_{mod} signal in the photocurrent. By mixing this beatnote down to DC, there is a non-zero voltage generated when the lock beam is near resonance. The full shape of this error signal is plotted in Fig. 2.9, highlighting the linearity of the signal in a region of width $\kappa/2$ near $\Delta=0$.

This PDH locking scheme forms the basis of multiple locks in our experiment. To stabilize the detuning between a laser and a cavity, one can apply feedback to either the




Figure 2.9: Sample PDH error signal.

laser or the cavity frequency (for example, via its length). We use a combination of both approaches, in order to achieve simultaneous locking of our nested feedback system. We implement proportional-integral (PI) feedback using a standard commercial ‘servo controller’ (New Focus LB1005), which allows for easy adjustment of offsets to the error signal and feedback output, as well as the feedback gain and PI bandwidth.

For stabilizing the lasers to the experimental cavity, a PDH error signal is generated from the measurement beam, and feedback is applied in two ways: First, feedback at frequencies up to 5kHz is applied to the frequency of the 200MHz AOM through which all the beams pass (see Fig. 2.6). The PI output signal is low-pass filtered and applied to the voltage-controlled oscillator (VCO) which drives the AOM. This AOM has a limited tuning range, so it is necessary to keep its frequency centered within that range by handling any long-term frequency drifts of the cavity/laser via some other feedback channel. We choose to apply slow feedback to the measurement laser piezo. Since the quantity we wish to stabilize is the VCO frequency, we simply split its driving voltage (the output of the main PI feedback loop), and use it as the input to a second PI controller (LB1005-2). The output of this controller is low-pass filtered around 1Hz, to ensure that we are only compensating for slow drifts.

As mentioned previously, any feedback applied to the 200MHz AOM will shift all of the beams simultaneously, which is useful, since all beams should share the same detuning fluctuations. However, for this to work, it is necessary that the control laser frequency be stabilized to the measurement laser frequency (both to eliminate intrinsic fluctuations

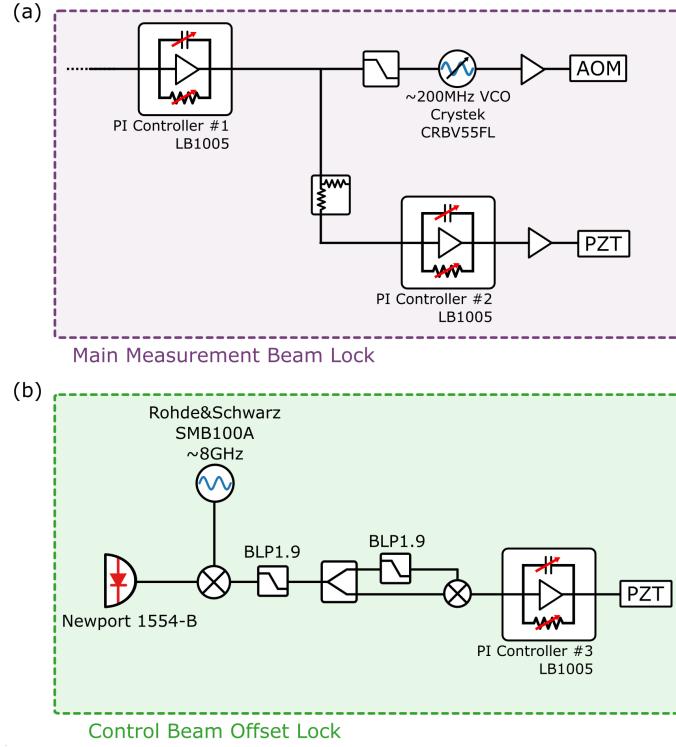


Figure 2.10: Feedback circuits for the main measurement beam lock and the control beam offset lock.

between the lasers, as well as duplicate the slow feedback being applied to the measurement laser piezo). To achieve this frequency lock, we begin by mixing small samples of each beam (taken immediately after the laser outputs) on a high-speed photodiode. The photocurrent contains a microwave signal at $\sim 8\text{GHz}$ (the beat note of the lasers), which we wish to stabilize in frequency. The circuit used to generate a usable error signal for this beat note is depicted in Fig. 2.10. The photocurrent is mixed with a reference signal at $\omega_{\text{ref}} \approx 2 \times \text{FSR}$, such that the beat note between the lasers is mixed down to a signal near DC. To generate a usable error signal, we need a voltage which is linearly proportional to the frequency of this beat note. To achieve this, we take the output of the first mixer, split it, send one arm through a 1.9MHz low-pass filter, then mix the two arms together again. Thus, whenever the frequency difference of the lasers is near $\omega_{\text{ref}} \pm 1.9\text{MHz}$, there will be a 1.9MHz signal in each arm, but one will be delayed by the filter, such that the mixer output is a voltage proportional to the original laser frequency difference. This error signal is fed into another PI controller (LB1005-3), which applies feedback to the cavity piezo of the control laser.

Finally, to keep both lasers on-resonance with their respective filter cavities, PDH feed-

back is used again, but now with the feedback applied to piezos which adjust the filter cavity lengths, keeping them resonant with the incident lasers. When locking the filter cavities to the lasers, it is particularly useful that the majority of the experimental cavity feedback is applied to the AOM located after the filter cavities. The experimental cavity has significant low frequency (\sim kHz) frequency noise, spanning hundreds of kHz in range; if the lasers were being modulated directly to follow these fluctuations, it would be necessary to have the filter cavities track this large noise, which would be difficult. Instead, the laser frequencies are mostly stationary, reducing the amount of feedback necessary to the filter cavities – in fact, one can tune the filter cavities to be approximately resonant with the laser for \sim seconds without any feedback.

Chapter 3

Exceptional Points: Degeneracies in Open Systems

The work presented in this thesis is largely focused on the idea of adiabatic, topological dynamics near an ‘exceptional point’ degeneracy. This physics can be thought of as an extension of traditional notions of degeneracies and adiabatic control to open systems (i.e. systems subject to loss/gain via coupling to their environment). Hence, we will begin with a review of degeneracies in closed systems, including the idea of adiabatic control of such systems near a degeneracy. Then, these ideas will be extended to an open (non-Hermitian) system, again making use of a minimal model system which illustrates the main behavior associated with EPs. We will first consider the structure of the eigenvalue surfaces near an EP, finding that they possess a non-trivial topology which underlies many of the effects we’ll study. In particular, we’ll explore how this topology influences the behavior of the system with respect to closed adiabatic trajectories. This discussion will reveal some of the core features of interest of EPs, namely energy transfer which is topology-dependent and non-reciprocal, due to the presence of loss and gain. We’ll finish by reviewing the state of the field, conceptually and experimentally, in order to establish the context of the measurements presented in the next chapter.

3.1 Closed Systems: Avoided Crossings and the Adiabatic Theorem

We begin by considering a generic, coupled, two-mode system, in which the frequencies of the modes (ω_1, ω_2) are tunable, as well as the coupling rate (g) . Both the frequencies and the coupling are assumed to be real. We'll keep this model intentionally system-agnostic, as it can be applied in a wide variety of scenarios. For instance, in discussions of adiabaticity in quantum mechanics, a common toy example is a spin-1/2 particle in an external magnetic field¹. But mathematically, this same model is equally valid for describing coupled classical oscillators, as discussed in Chs. 1 and 4. So, for now, we'll just suppose that we have a 2-mode system governed by the following equation of motion:

$$\dot{\vec{c}} = -i\mathbf{H}\vec{c} \quad (3.1)$$

where

$$\mathbf{H} = \begin{pmatrix} \omega_1 & g \\ g & \omega_2 \end{pmatrix} \quad \text{and} \quad \vec{c} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \quad (3.2)$$

As we begin to consider an eigenvalue analysis of this system, the mean frequency $\bar{\omega}$ will simply apply an overall shift to the eigenvalues. Since this does not qualitatively affect the dynamics, we'll move to a frame rotating at $\bar{\omega}$, leaving a traceless ($\text{tr}(\mathbf{H}) = 0$) Hamiltonian described in terms of the frequency splitting $\Delta = (\omega_1 - \omega_2)/2$.

$$\mathbf{H} = \begin{pmatrix} \Delta & g \\ g & -\Delta \end{pmatrix} \quad (3.3)$$

The eigenvalues of this simple system are given by $\lambda_{\pm} = \pm\sqrt{\Delta^2 + g^2}$. A common exercise is to suppose that the frequency splitting of the modes is tunable (with a fixed, non-zero coupling), and ask how the eigenvalues of the system behave. The result, shown

1. Note, however, that in the following discussion, we'll be considering a real matrix, which only describes a spin in a two-dimensional magnetic field.

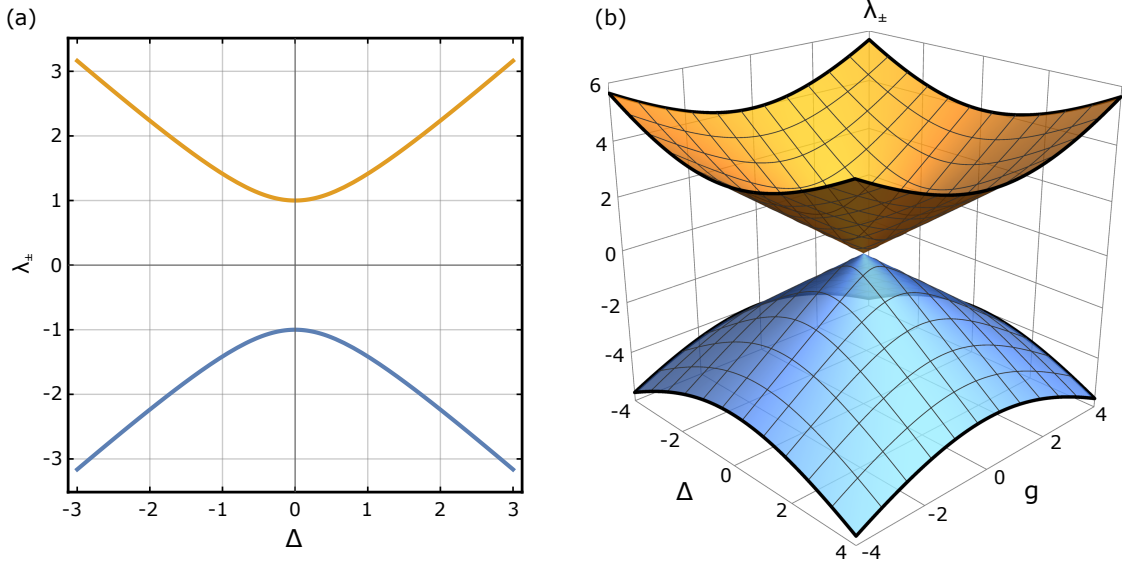


Figure 3.1: Eigenvalues of a closed system. (a) Eigenvalues as a function of Δ for $g = 1$, showing an avoided crossing with gap size given by twice the coupling. (b) Eigenvalues as functions of g and Δ , showing cones emanating from a degeneracy (DP)

in Fig. 3.1a for $g = 1$, is that the modes experience an avoided crossing² as Δ is swept through zero. It is this scenario (i.e. the Landau-Zener problem) in which one often first encounters the adiabatic theorem [15], which we'll review briefly.

Suppose that we apply a time-dependent perturbation to the system, by tuning the parameter Δ . Let the instantaneous eigenmodes and eigenvalues of the system be given by $\psi_m(t)$ and $\lambda_n(t)$. Then, if the system was prepared in a particular eigenmode at time $t = 0$ ($\Psi(0) = \psi_n(0)$) and if the perturbation is applied sufficiently slowly, then the time dependence of the state vector will be given by:

$$\Psi(t) = \psi_n(t) e^{i\theta_n(t)} e^{i\gamma_n(t)} \quad (3.4)$$

where $\theta_n(t) = -\int_0^t \lambda_n(t') dt'$ is the ‘dynamical’ phase (simply the accumulation of phase due to the time-dependent frequency) and $\gamma_n(t) = i \int_0^t \langle \psi_n(t') | \dot{\psi}_n(t') \rangle dt'$ is called the ‘geometric’ phase (arising from the changing eigenbasis along the trajectory). Physically, the implication of this result is that if we perturb the system sufficiently slowly, then it will remain in the same instantaneous eigenmode throughout, simply accumulating a couple of

². Also called an ‘anti-crossing’ or ‘level repulsion’

phase factors. Thus, if we prepare our system in, for example, the upper branch of Fig. 3.1(a) and slowly vary Δ , it will remain in the upper branch throughout the process.

Holding on to this point for the moment, let's return to the eigenvalues of our two-mode system, and suppose that we have control of both Δ and g as tunable parameters. In this case, we can visualize the eigenvalues as two surfaces, seen in Fig. 3.1(b). We find two cones, at whose intersection lies a degeneracy point³ (DP). The avoided crossing we saw previously is essentially one slice of these surfaces parallel to the Δ axis. (Now it's clear that we would see the same for a slice parallel to the g axis.)

The adiabatic theorem would now predict that we can vary both g and Δ slowly, and expect the system to remain on either the upper or lower cone. (Note that full proof of the adiabatic theorem requires non-degeneracy, so it is important that we not pass through the DP.) A particularly interesting feature, noted by Michael Berry [16] (and some years before by Shivaramakrishnan Pancharatnam [17]) relates to the result of the adiabatic theorem when applied to closed loops in the parameter space around a DP. In such a scenario, the geometric phase can be written as a surface integral of a quantity called the Berry curvature. The value of this integral is determined by the presence of degeneracies. Curiously, this phase does not depend on the exact time-dependence of one's trajectory in parameter space – only the geometry of the trajectory determines the phase factor. In the particular case we've presented here (a real, symmetric matrix), it turns out that the DP has co-dimension two⁴, and the geometric phase integrates to either $\pm 2\pi$ or 0, depending on whether the curve encloses the DP, and in which sense. (In this sense, the geometric phase here becomes a topological phase: local deformations of the curve do not change the phase, only inclusion/exclusion of the degeneracy matters).

While this was a fairly cursory overview of adiabaticity and Berry phase, we can see how eigenvalue topology and the presence of degeneracies can have important global effects on closed adiabatic circuits. We'll now seek to build upon these ideas, and see how they change upon moving to describe open systems.

3. Sometimes called a 'diabolical point', as the intersecting cone shape is known as a 'diabolo'

4. Put quite briefly, this means that the space of Hamiltonians we're considering is spanned by two parameters, so we can unambiguously enclose the degeneracy with loops in this 2D control space.

3.2 Open Systems & Complex Eigenvalues

‘Open systems’ are systems which experience either loss or gain due to coupling with their environment. In a closed system, we considered eigenmodes with frequency ω_n , whose time-dependence was given by oscillations $e^{-i\omega_n t}$ of the normal coordinate. Now, we will now describe open modes via complex eigenvalues $\lambda_n = \omega_n + i\gamma_n$ whose time-dependence is both oscillatory and real-exponential: $e^{-i\lambda_n t} = e^{-i\omega_n t} e^{\gamma_n t}$.

3.2.1 Coupled Modes in an Open System

Let’s begin by considering the same two-mode system as before, but with each mode possessing a gain/loss⁵ rate γ_i :

$$\begin{pmatrix} \dot{c}_1 \\ \dot{c}_2 \end{pmatrix} = -i\mathbf{H} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = -i \begin{pmatrix} \omega_1 + i\gamma_1 & g \\ g & \omega_2 + i\gamma_2 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \quad (3.5)$$

As before, we note that from this effective Hamiltonian matrix, we can remove a diagonal matrix containing the mean frequency $\bar{\omega} = (\omega_1 + \omega_2)/2$ and mean decay rate $\bar{\gamma} = (\gamma_1 + \gamma_2)/2$, leaving a matrix defined in terms of the frequency splitting $\Delta = (\omega_1 - \omega_2)/2$ and relative gain/loss $\gamma = (\gamma_1 - \gamma_2)/2$:

$$\mathbf{H} = \mathbf{H}_0 + \mathbf{H}' = \begin{pmatrix} \bar{\omega} + i\bar{\gamma} & 0 \\ 0 & \bar{\omega} + i\bar{\gamma} \end{pmatrix} + \begin{pmatrix} \Delta + i\gamma & g \\ g & -\Delta - i\gamma \end{pmatrix} \quad (3.6)$$

The benefit to this separation is that the diagonal matrix \mathbf{H}_0 will simply offset the eigenvalues we are interested in⁶, without qualitatively affecting the dynamics. This separation can be thought of as ‘removing the trace’, since the matrix \mathbf{H}' we are left with is traceless. Going forward, we’ll simply redefine $\mathbf{H} \rightarrow \mathbf{H}'$, ignoring the trivial effects of including \mathbf{H}_0 . The eigenvalues of this new matrix \mathbf{H} are given by

5. $\gamma_i > 0$ corresponds to gain, $\gamma_i < 0$ corresponds to loss

6. If $\{\lambda_i\}$ are the eigenvalues of a matrix \mathbf{M} and \mathbf{D} is a diagonal matrix with entries μ_i , then the eigenvalues of $\mathbf{M} + \mathbf{D}$ are $\{\lambda_i + \mu_i\}$.




Figure 3.2: Eigenvalues of \mathbf{H} as a function of Δ , for various values of g . The eigenvalues are split into their real and imaginary parts. While coupling in a closed system always led to an avoided crossing (except at the DP), here we find that the crossing can occur in either the real or imaginary parts, depending on the value of g .

$$\lambda_{\pm} = \pm \sqrt{g^2 + (\Delta + i\gamma)^2} \quad (3.7)$$

By virtue of considering eigenvalues which are complex numbers, there is inherently more information to consider when visually representing these eigenvalues. Following the same approach as in Sec. 3.1, we might consider plotting the eigenvalues as one of the system parameters (for instance, Δ) is varied. This is plotted in Fig. 3.2, with the real and imaginary parts of λ_{\pm} split into separate axes. Here, we see that for some values of g (say, $g = 1.25$), the mode frequencies ($\text{Re}[\lambda_{\pm}]$) experience an avoided crossing just as in the closed system, while the imaginary parts ($\text{Im}[\lambda_{\pm}]$) cross. Meanwhile, other values of g (say, $g = 0.75$) result in a crossing of frequencies, and an anti-crossing of imaginary components. This is already distinct from the closed-system case, in which the mode frequencies always experienced an avoided crossing as g or Δ were varied. Fig. 3.2 also shows the eigenvalues when $g = 1.0$, in which we see both the real and imaginary parts crossing, with a degeneracy in both parts at $\Delta = 0$. Note also the eigenvalues have a sharper (square-root) dependence at this degeneracy, compared to the linear crossing at a DP.

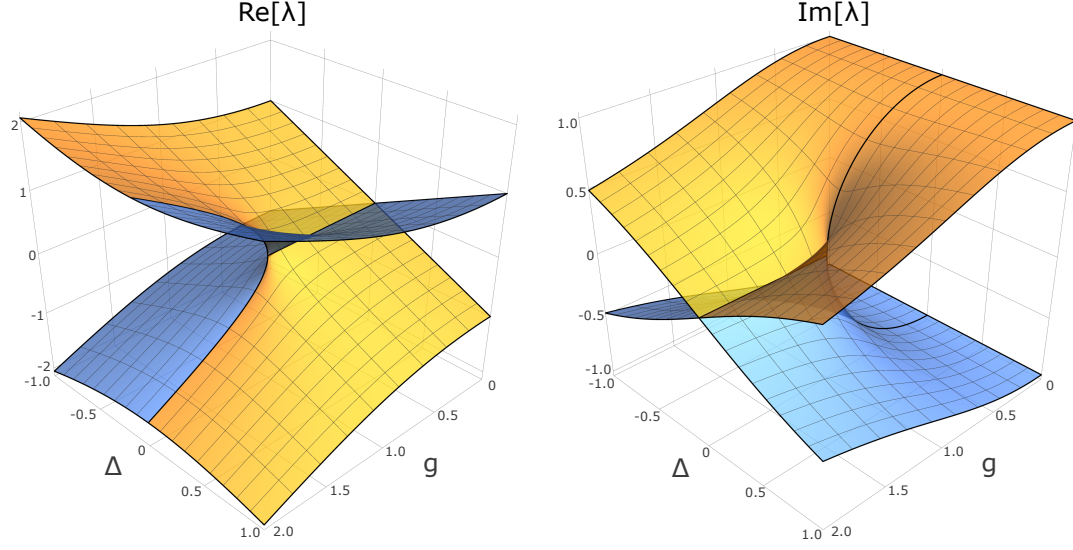


Figure 3.3: Eigenvalues of the coupled open system defined by \mathbf{H} , separated into real/imaginary parts and plotted as functions of g and Δ .

We'll now consider these eigenvalues in the full 2D control space formed by g and Δ , as we did in the closed system, but plotting separately the real and imaginary components of the eigenvalues. The resulting eigenvalue manifolds are shown in Fig. 3.3 (again, with $\gamma = 1$). At the center of the sheets lies the point $(g, \Delta) = (1, 0)$, at which both the real and imaginary parts of the eigenvalues coalesce. This point is the aforementioned 'exceptional point', which lies at the heart of this work. Among many interesting mathematical properties of this point, the central feature of interest is the topology of the eigenvalues around the EP. Specifically, the EP is a branch point degeneracy with non-trivial monodromy, meaning that analytic continuation around the EP results in a swap of eigenvalues. This is the sort of behavior one encounters frequently in complex analysis, for example when considering the complex square-root function \sqrt{z} or complex logarithm $\log(z)$. (In fact, the sheets in Fig. 3.3 are

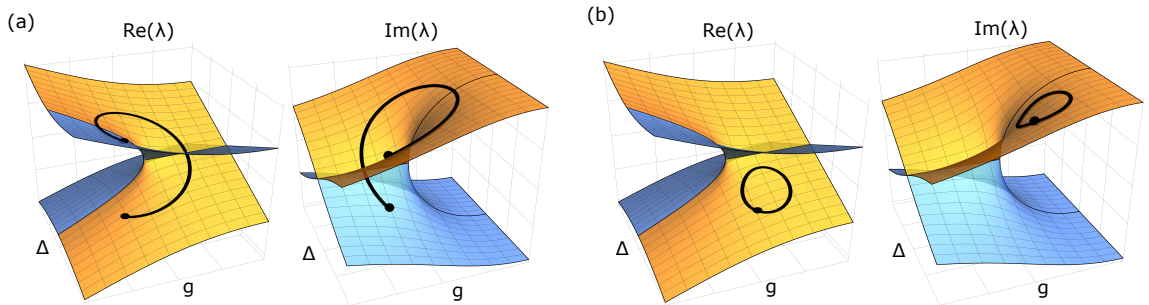


Figure 3.4: Complex eigenvalue sheets with trajectories which do (a) and do not (b) enclose the EP.

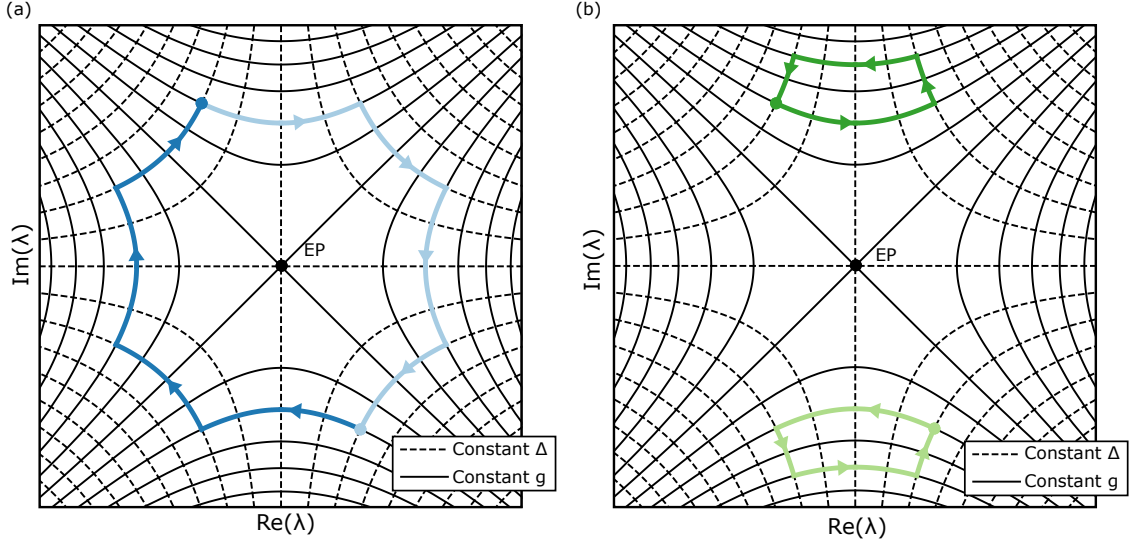


Figure 3.5: Complex eigenvalues contours and trajectories. Both (a) and (b) show the complex eigenvalues of \mathbf{H} along contours of constant g (solid lines) and Δ (dashed lines). For any particular (g, Δ) combination, there are two eigenvalues, such as illustrated by the pair of blue dots in (a) and the pair of green dots in (b). The blue and green curves show the smooth trajectories of these eigenvalues as g and Δ are varied cyclically along a closed set of contours which do and do not enclose the EP, respectively.

topologically equivalent to the complex square-root function, since the eigenvalues arise from calculating a square root.) This monodromy is qualitatively illustrated by the black curve in Fig. 3.4(a), which shows a trajectory which encloses the EP, beginning with one eigenvalue and ending with the other. For comparison, a loop which does not enclose the EP is also plotted (Fig. 3.4(b)), clearly returning to its original value. This exchange of eigenvalues under encirclement of the EP is one of the key features which has drawn interest to these systems in recent decades.

3.2.2 Alternative Perspective: Complex Avoided Crossings

One might find some additional intuition for these systems by considering their eigenvalues presented in a different fashion. In Fig. 3.5, we plot our two eigenvalues in the complex plane, along contours of constant g and Δ , with the EP marked in the center. One sees here that (except for two paths which cross the EP), one can still think of the system as experiencing avoided crossings, if one considers these crossings in the complex plane. Here, we see that while sometimes the real parts cross, and other times the imaginary parts cross, there is always a repulsion in λ_{\pm} , with a gap size set by the coupling strength g .

In Fig. 3.4, we saw how one eigenvalue can be analytically continued to the other, by encircling the EP. The same behavior is illustrated in Fig. 3.5(a). The dark and light curves present the trajectory of both eigenvalues as g and Δ are varied along a closed path, following contours of constant g and Δ , alternately. The two blue dots indicate the eigenvalues at the beginning/end of the trajectory. Since the path in (g, Δ) space encloses the EP, the two eigenvalues switch, beginning at one dot and ending at the other. This is illustrative of the same monodromy we saw before. Fig. 3.5(b) illustrates the eigenvalue trajectories along a closed path that does not enclose the EP, and hence does not exchange eigenvalues. We see each trajectory returning to the same eigenvalue it started at, as in Fig. 3.4(b).

3.3 Adiabatic Control in Open Systems

Having illustrated the topology of eigenvalues near an EP, we might now ask what impact this has on adiabatic control of the system. Recall that in closed systems, we stated that sufficiently slow variation of the Hamiltonian would result in a system remaining in the instantaneous eigenstate, as long as the modes remained non-degenerate. Given the monodromy near an EP, it now seems that closed adiabatic perturbations could actually cause the system to evolve from one eigenmode to another. If so, this would allow adiabatic energy transfer based solely on a topological property of the control loop (i.e. encirclement of the EP). In the next few sections, we'll see next is that this sort of energy transfer is indeed possible, for certain control loops and initial conditions. However, we will also find that the presence of loss and gain in the system introduces non-trivial differences that restrict a naïve application of the adiabatic theorem.

In simple closed systems (like the one from the previous section), analytic solutions for the result of a time-dependent perturbation allow one to calculate state evolution explicitly (i.e. the Landau Zener formula). In open systems, some analytic results exist for toy models like the one presented here [18–20], allowing one to calculate, for example, the adiabatic (or non-adiabatic) behavior of the system. This is possible because in certain regimes, one can find analytic expressions for the eigenbasis of this toy model, and calculate state propagation

under simple perturbations of the control parameters (linear, or circular paths, for example). Outside of these approximation regimes, or for any perturbation which does not match this simple model, analytic solutions will be harder to find. In later chapters, for example, we'll work with an experimental realization of an EP which is topologically equivalent to this toy model, but which has eigenvalues with a very complicated dependence on the control parameters. So, instead of seeking analytic results, our general approach to understanding these systems will be numerically integrating the equations of motion, given some particular time-dependent control parameters. We begin with the toy model introduced previously:

$$\dot{\vec{c}} = -i\mathbf{H}\vec{c} = \begin{pmatrix} \Delta + i\gamma & g \\ g & -\Delta - i\gamma \end{pmatrix} \vec{c} \quad (3.8)$$

To consider adiabatic perturbations, we'll allow for time-dependent control parameters $\Delta(t)$ and $g(t)$, choosing circular paths in the $\{\Delta, g\}$ control space and executed uniformly in a time τ . We'll describe these loops by their center (Δ_0, g_0) , radius (r_0) , and a phase offset (ϕ_0) which controls where along the circle the loop begins/ends:

$$\Delta = \Delta_0 + r_0 \sin\left(2\pi\left(\frac{t}{\tau} + \phi_0\right)\right) \quad (3.9)$$

$$g = g_0 + r_0 \cos\left(2\pi\left(\frac{t}{\tau} + \phi_0\right)\right) \quad (3.10)$$

Note that these expressions describe a clockwise loop. Counter-clockwise loops can be described by simply changing the sign of t in the above equations.

To visualize state evolution during some applied perturbation, we will need the projection onto the instantaneous eigenbasis. Let $\vec{v}_{\pm}(t)$ be the instantaneous eigenvectors of $\mathbf{H}(t)$. For all times, the solution $\vec{c}(t)$ can be written as a projection onto this basis:

$$\vec{c}(t) = c_{-}(t)\vec{v}_{-}(t) + c_{+}(t)\vec{v}_{+}(t) \quad (3.11)$$

One naïve (and unjustified) approach to thinking about this system would be to assume that adiabaticity (in the sense of smooth evolution along the eigenvalue surface) would

occur as long as the perturbation is slow compared to the energy gap ($|\lambda_+ - \lambda_-|$). That is, if we were to prepare $c_-(0) = 1$ and $c_+(0) = 0$, then for all time t , $c_-(t)$ would remain the dominant projection⁷. To see if this occurs, we can simply numerically integrate the linear differential equations of Eq. 3.8, using standard computational techniques⁸. This calculation yields a full solution for the time dependence of the system, $c(\vec{t})$. One should note, however, that this result is in the basis of the differential equation, not the eigenbasis (i.e. $c_{1,2}(t)$, not $c_{\pm}(t)$). To find the projection into the instantaneous eigenbasis, one should apply a change of basis defined by a matrix of the instantaneous eigenvectors:

$$\begin{pmatrix} c_+(t) \\ c_-(t) \end{pmatrix} = \mathbf{N}(t) \begin{pmatrix} c_1(t) \\ c_2(t) \end{pmatrix} \quad (3.12)$$

where

$$\mathbf{N}(t) = \left[\begin{pmatrix} \vec{\nu}_+(t) & \vec{\nu}_-(t) \end{pmatrix}^T \right]^{-1} \quad (3.13)$$

There are various ways to visualize these projections, in order to understand the behavior of the system. One way to construct a particularly useful quantity⁹ is to compute a weighted average of the eigenvalues:

$$\bar{\lambda}(t) = \frac{|c_+(t)|^2 \lambda_+(t) + |c_-(t)|^2 \lambda_-(t)}{|c_+(t)|^2 + |c_-(t)|^2} \quad (3.14)$$

Thus, if the state predominantly projects along $\vec{\nu}_+$ ($\vec{\nu}_-$), then $\bar{\lambda} \approx \lambda_+$ (λ_-). For intermediate values, $\bar{\lambda}$ will lie somewhere between the two eigenvalues. The normalization by $|c_+(t)|^2 + |c_-(t)|^2$ is important, since the presence of gain and loss means that the total energy of the two modes will not be conserved.

Fig. 3.6 shows the result of simulating the system under perturbation by a loop with

7. Note that one must be careful about working with a continuous eigenbasis here – As we’ve already seen, our eigenmodes can be analytically continued into each other, so clearly there will be some point at which the \pm naming convention changes. This is a detail which is important for analytical calculations, but will not significantly affect any of our work here.

8. For these numerical simulations, I use Mathematica’s `NDSolve` function.

9. One practical advantage of this quantity: It is symmetric with respect to \pm . Thus, the problem of working with a continuous eigenbasis is avoided – Even if the \pm assignment changes during one of our calculations, $\bar{\lambda}$ will not change.

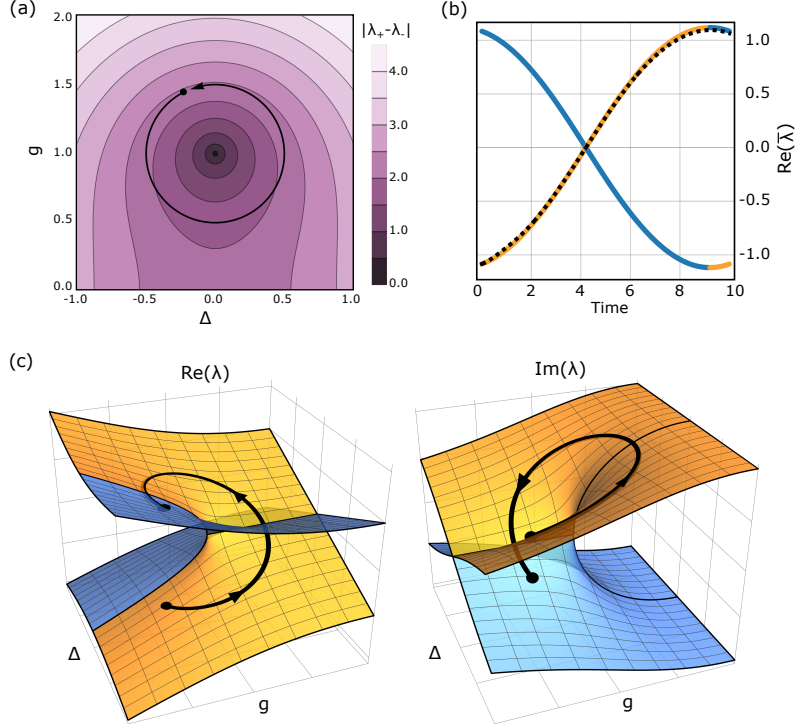


Figure 3.6: System trajectory under perturbation. (a) Control loop, overlaid on a contour plot of the eigenvalue difference (to illustrate the location of the EP) (b) Real projection of $\bar{\lambda}$, along with the real eigenvalues throughout the control loop. The color of the eigenvalue indicates whether it has the larger (orange) or smaller (blue) imaginary part (the same convention used in the sheets). (c) Projection of $\bar{\lambda}$ on the real and imaginary sheets.

$(r_0, g_0, \Delta_0, \phi_0) = (0.5, 1.0, 0.0, 0.1)$. The control loop is illustrated in Fig. 3.6(a), and in Fig. 3.6(c), we see $\bar{\lambda}(t)$ separated into its real and imaginary parts, plotted on top of the corresponding eigenvalue sheets. Fig. 3.6(b) is the same result, reduced to a two-dimensional plot, showing only the real part of $\bar{\lambda}$ (black dashed curve), as well as the real parts of the eigenvalues λ_{\pm} along the control loop. The color of the eigenvalues indicates which one has the greater imaginary part (just as it does in the sheets). In both views, we see that for this particular time-dependent perturbation, the state vector seems to follow the instantaneous eigenmode, consistent with what one might expect from the adiabatic theorem.

The behavior demonstrated in Fig. 3.6 is our first example of the aforementioned ‘energy transfer’, to which much of this thesis will be devoted. We have mentioned before that this energy transfer is the result of some adiabatic-like process, and should depend on the geometry of the loop, in relation to the topology of the eigenvalue manifolds. These

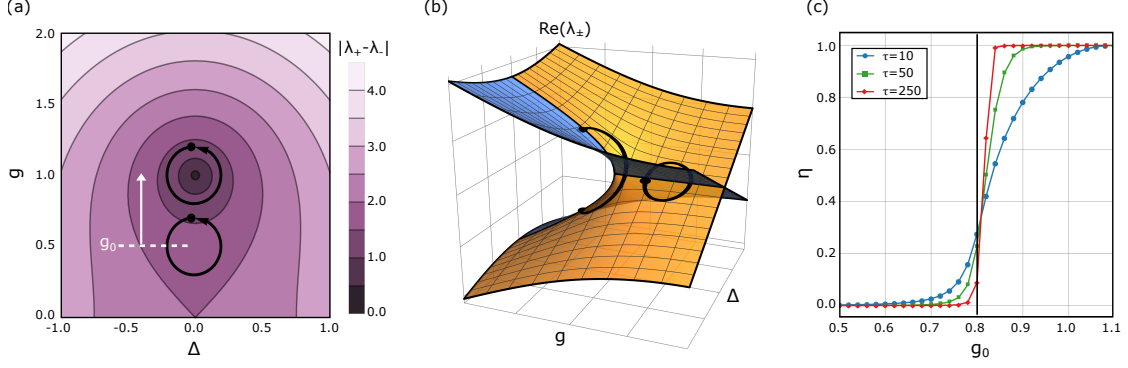


Figure 3.7: Varying control loop location. (a) Illustration of control loop, relative to EP. Colored contour regions show the eigenvalue difference, which vanishes for $(g, \Delta) = (1.0, 0.0)$. The loop is centered at $(g = g_0, \Delta = 0.0)$, and g_0 is varied from 0.5 to 1.1, as indicated. (b) Sample simulated trajectories (shown on the real eigenvalue sheets) for the two control loops shown in (a). One encloses the EP and result in energy transfer; the other does not enclose the EP, and returns to its starting point. (c) Results of systematically varying the loop location, for different values of τ . The black line indicates the value of g_0 at which the loop begins to include the EP.

details are what we'll now demonstrate via this toy model, using simulations like the one just presented. In order to systematically characterize this behavior, it will be useful to have a metric for quantifying energy transfer. Essentially, if we initialize the system into a particular mode ($c_{\pm}(0) = 1$), we would like to know the amplitude of the other mode at the end of the loop ($c_{\mp}(\tau)$). However, due to growth/decay of the overall system energy, we should also normalize by the total energy at the end of the loop. Thus, we'll define a quantity called the 'relative transfer efficiency', η :

$$\eta = \frac{|c_{\pm}(\tau)|^2}{|c_{+}(\tau)|^2 + |c_{-}(\tau)|^2} \quad (3.15)$$

where c_{\pm} indicates the mode which is not initialized. So, for example, if we prepare the system in the c_{+} mode at $t = 0$, then η describes how much energy has been transferred to the c_{-} mode at $t = \tau$ (relative to the total energy at this time). So, for a c_{+} initialization, if c_{-} is the dominant projection at the end, then we find $\eta \sim 1$, whereas if the dominant projection is still c_{+} , then we have $\eta \sim 0$.

3.3.1 Dependence on Eigenvalue Topology

With this metric, we'll now consider how the energy transfer illustrated in Fig. 3.6 depends on the loop geometry (in relation to the eigenvalue topology). Our intuition for

this process indicates that the energy transfer should only occur when the loop in control space encloses the degeneracy. In Fig. 3.7, we confirm this systematically, by translating a control loop along the g -axis, as depicted in Fig. 3.7(a). The loop is defined by $(r_0, \phi_0, \Delta_0) = (0.2, 0.01, 0.0)$, and with a center, g_0 , which is varied from 0.5, to 1.1. The value $g_0 = 0.8$ corresponds to the boundary where the loop transitions from not enclosing the EP to enclosing the EP. Fig. 3.7(c) shows how the relative transfer efficiency varies with the loop location (described by the loop center g_0), and indeed we see the energy transfer transition from ~ 0 to ~ 1 once the loop begins enclosing the EP. This behavior is shown for several different loop durations: $\tau = 10, 50, 250$. We will conduct a more systematic study of adiabaticity in the next section, but already we can understand this time-dependence, based on our intuition from the closed-system adiabatic theorem. Consider first the faster loops ($\tau = 10$). For values of g_0 near 0.8, the control loop is passing very near to the EP, which means that the difference in the eigenvalues is becoming smaller (since it vanishes at the EP). Recalling that in a closed system, the condition for adiabaticity is set by the gap in eigenvalues, this suggests that a slower loop speed is needed for these loops which pass near the EP. Indeed, we see that if we increase the loop time (to $\tau = 50$ and $\tau = 250$), we find that the onset of energy transfer occurs more sharply at $g_0 = 0.8$. This gap condition is of course not the full story for predicting adiabatic behavior in this system, as we've suggested already, but it provides reasonable intuition in some situations, including the example shown here.

3.3.2 Dependence on Perturbation Speed

Let's now return to the large, EP-enclosing control loop from Fig. 3.6 and consider how this energy transfer depends on the timescale of our perturbation. We've already seen that a loop time of $\tau = 10$ was apparently slow enough for the state to follow the instantaneous eigenmode during the perturbation. As an example which is distinctly outside this regime, consider the same control loop, applied during a time $\tau = 0.1$. This sudden perturbation corresponds to a 'diabatic' process, and we might expect the state vector to remain essentially unchanged, occupying the same eigenmode it did at the beginning of the control loop. Indeed, Fig. 3.8 shows that the state trajectory returns to its initial eigenmode.




Figure 3.8: Calculated trajectory for a diabatic (sudden) perturbation. The system essentially remains in its initial state, and the black curve in (a) corresponds to the result of projecting this state onto the instantaneous eigenbasis. (b) shows the real part of $\bar{\lambda}$, as it compares to the real parts of λ_+ and λ_- (blue and orange). Also shown (in green) is the trajectory we would see if the state remained exactly in the initial state.

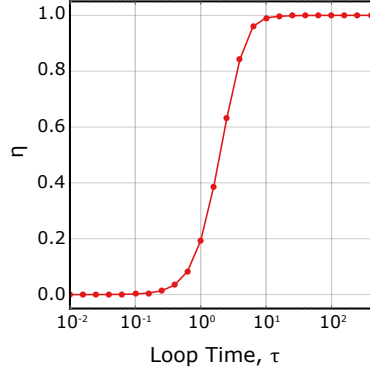


Figure 3.9: Varying control loop duration. Relative transfer efficiency η , as a function of loop time τ , showing the onset of adiabatic energy transfer around $\tau \approx 2$.

The trajectory it follows is simply the result of projecting the initial state vector onto the instantaneous eigenbasis. To emphasize this, a green curve has been added to Fig. 3.8(c), which shows the mean frequency ($\text{Re}(\bar{\lambda})$) of the initial eigenvector, projected onto the instantaneous eigenbasis (i.e. the path that we would expect if the state vector was equal to the initial eigenvector for all times). We see that for short times like this, the state vector does not have time to evolve significantly away from its initial orientation.

Having considered examples from both extremes, we could now use our previously defined transfer efficiency η to systematically study this dependence on adiabaticity. Fig. 3.9 illustrates how η depends on the loop time τ , with an apparent onset of adiabaticity around $\tau = 2$.

Thus far, we’ve identified the ways in which this energy transfer is consistent with our expectations based on the topology of the eigenvalues and a naïve application of the adiabatic theorem. However, such expectations do break down in open systems, specifically because of the relative gain and loss, as we shall see next.

3.4 Non-adiabatic behavior

Recall the example of adiabatic energy transfer introduced in Fig. 3.6, which began in the lower-frequency eigenmode and ended in the higher-frequency eigenmode, after application of a control loop which was counter-clockwise in the control space. The control loop duration was $\tau = 10$, which in the previous section we found was sufficiently long to allow adiabatic evolution. In Fig. 3.10, we show the result of simulating the same perturbation, but now when the system is initialized in the higher-frequency eigenmode. The behavior is distinctly different, as the state trajectory no longer follows the instantaneous eigenmodes, instead ‘jumping’ from one eigenmode manifold to the other during the first half of the loop. As a result of this jump between eigenmodes, the system returns to the original eigenmode once the perturbation is complete. Since the eigenvalue gap along this path is the same as before, there is no reason to suspect that we are failing to satisfy the adiabaticity condition. Indeed, systematically varying the loop duration shows that Fig. 3.10 is already the long-time behavior of the system. It would seem that adiabatic evolution along this particular path is not possible – revealing an apparent gap in the story presented thus far. As we’ll see, the key distinction between this trajectory and the previous one is the relative gain/loss of the instantaneous eigenmodes during the perturbation.

Recall that in choosing how to color-code the two eigenvalue manifolds, we have always chosen to distinguish the two modes based on the sign of their relative loss/gain. That is, the orange eigenmode is always the one with a positive imaginary part (gain), and the blue is the one with a negative imaginary part (loss)¹⁰. In the first example we considered (the one which was found to adiabatically follow the eigenmodes), the trajectory was one

10. In this particular model, one imaginary part is always positive and the other is always negative. In a more generic scenario, what matters for this discussion is which imaginary part is greater than the other (even if both are positive or negative)




Figure 3.10: Calculated trajectory for the same perturbation as Fig. 3.6, but now beginning in the higher-frequency mode. The first part of the trajectory attempts to follow the loss (blue) sheet, and therefore experiences a non-adiabatic jump to the orange sheet. On the right, we see the trajectory’s real component, on top of the real parts of the eigenvalues. As before, the coloring indicates gain/loss, and see that the trajectory jumps from the blue (lossy) mode to the orange (gain) mode near the beginning

which lay primarily on the orange (gain) eigenmode. In contrast, when we initialized in the other eigenmode, the adiabatic trajectory would have primarily followed the blue (lossy) eigenmode. As it turns out, gain/loss asymmetry makes extended propagation along the lossy mode impossible [20].

To understand this, we’ll recall the idea underlying adiabaticity in closed systems: If the system is prepared in one eigenmode, then a time-dependent perturbation of the Hamiltonian may result in a small projection onto the opposite eigenmode, but this non-adiabatic ‘leak’ can be suppressed by applying the perturbation more slowly. Suppression of this leaked projection is what fails in an open system. If we suppose that our system is prepared in the lossy eigenmode, and we begin perturbing the Hamiltonian, as soon as the state vector develops a projection into the gain mode, that component will begin to grow exponentially, due to the gain. As this projection into the gain mode grows, it eventually dominates the mode decomposition, reflecting an effective transition from the loss eigenmode to the gain eigenmode. If one applies the perturbation more slowly, the ‘leak’ will be exponentially suppressed (as it is in the closed system), but will also have more time to grow from the gain, leading to the same result. Thus, a slower perturbation does not avoid the transition to the gain mode, and in fact this non-adiabatic behavior is unavoidable.

This non-adiabatic behavior due to growth in the gain mode occurs anytime the system attempts to propagate along the loss mode. But, if the time spent in the loss mode is

short compared to the time spent in the gain mode, the system can effectively propagate adiabatically, since the ‘leak’ into the gain mode will be insufficient to overwhelm the large amplitude entering into the loss mode. As noted previously, in the examples discussed thus far, the adiabatic trajectories lay almost entirely on either the gain mode or the loss. So, for example, even though the example from Fig. 3.6 did end in the loss mode, it spent only a small portion of the trajectory there, after spending most of the loop in the gain mode, and therefore we did not see any non-adiabatic jump¹¹ In the next example, the adiabatic trajectory would have spent most of the loop in the loss mode, and therefore there was ample time for a non-adiabatic jump to occur. Once the system jumps, it is able to adiabatically follow the gain mode back to where it began.

Consider the two examples we just presented, in which we applied the same perturbation to a system prepared in either of the two possible initial states. Due in one case to adiabatic transfer, and in the other to a non-adiabatic jump, we found that the result of the operation was a system occupying the higher-frequency mode. This behavior might be summarized by describing the entire perturbation in terms of a propagation matrix, mapping the initial state vector $\vec{c}(0)$ to the final state vector $\vec{c}(\tau)$:

$$\vec{c}(t) = \mathbf{U}\vec{c}(0) \quad (3.16)$$

In this notation, the behavior we just presented corresponds to the fact that \mathbf{U} must be asymmetric. This ability to affect asymmetric population transfer has been one of the main features drawing interests to EP systems. Fig. 3.11(a) illustrates the asymmetry of the propagation matrix for the perturbation presented in the previous two examples. If one were to interpret this propagation matrix as a scattering matrix, where the inputs are the modes at $t = 0$ and the outputs are the modes at $t = \tau$, then this behavior satisfies the usual definition of non-reciprocity (i.e. that $\mathbf{U} \neq \mathbf{U}^T$).

Having identified this asymmetric behavior, it can be interesting to consider what hap-

11. In Appendix A, we’ll refine this statement, and find that in an even longer time regime, there actually will be non-adiabatic transitions from the loss to gain mode, regardless of how little time is spent there. This is relevant for time and parameter regimes that are generally beyond the work considered here, so, for the moment, we will hold on to this intuition for predicting non-adiabatic jumps in the usual adiabatic time regime.




Figure 3.11: Calculated asymmetry of propagation matrix for clockwise and counter-clockwise perturbations. (Only the magnitudes of the complex matrix elements are plotted). In the diabatic time regime ($\tau \sim .01$), the final state is approximately equal to the initial state, and so we see that each propagation matrix is nearly the identity matrix. As we reach the adiabatic regime ($\tau \sim 10$), we expect to see an asymmetric propagator matrix, with the asymmetry switching between clockwise/counterclockwise. Indeed, for the clockwise perturbation we see $U_{12} > U_{21}$, while the counter-clockwise perturbation has $U_{12} < U_{21}$. One might also notice that the clockwise/counterclockwise matrices are equal to each other under transpose. This is consistent with the idea that these loops are time-reversals of each other.

pens if we reverse the perturbation from the previous two examples. For one thing, it is straightforward to see that if we consider a particular perturbation, then reverse the direction and initial condition, the adiabatic trajectory would essentially retrace itself. So, if the system actually followed such an adiabatic trajectory (as it did in the first example), then we would expect that reversing the direction and initial condition would also result in an adiabatic evolution. This behavior is summarized in Fig. 3.12, which compares simulated clockwise trajectories for both possible initial conditions. (These trajectories result from a perturbation of the same speed as the previous examples ($\tau = 10$) which was chosen to be in the ‘adiabatic’ regime¹² for all examples.)

Here, we note an interesting result: While the counter-clockwise perturbation resulted in preferential population of the higher-frequency mode (via adiabatic or non-adiabatic processes), we see that the clockwise perturbation results in preferential population of the lower-frequency mode. Thus, not only does adiabatic control near an EP allow asymmetric energy distribution, but the direction of that asymmetry can be chosen via the direction of the perturbation. This is one of the more powerful and interesting features which has

12. By this, we mean that at least one of the initial conditions behaves adiabatically. Or, in general, this is the time regime in which any transitory, dynamic effects are not important – only adiabatic evolution (or brief non-adiabatic jumps) govern the behavior of the system

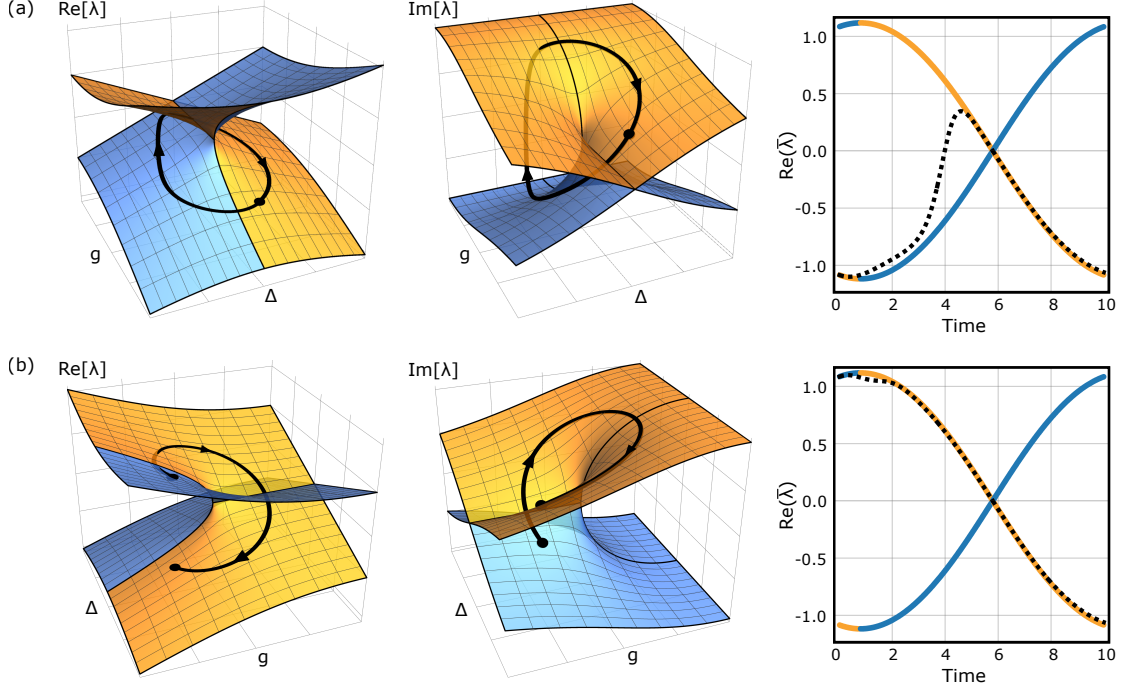


Figure 3.12: Trajectory simulations in which a clockwise perturbation is applied to a system initialized in either the lower-frequency mode (a) or the higher-frequency mode (b). In (a), the trajectory lies primarily on the orange (gain) manifold, and therefore we see adiabatic evolution and energy transfer from one mode to the other. In (b), the trajectory attempts to follow the loss manifold, resulting in a non-adiabatic jump to the gain manifold. As a result, the trajectory ends in the same eigenmode in which it began, and no energy transfer occurs.))

drawn interest to EP-based adiabatic control schemes. The propagation matrix elements for this clockwise perturbation are shown in Fig. 3.11(b).

For completeness, and to confirm that both perturbations reach our ‘adiabatic’ regime at the same time, we plot in Fig. 3.13 the energy transfer results for all four combinations of perturbation direction and initial condition. Figs. 3.11 and 3.13 are useful ways to visualize/summarize the behavior we’ve just discussed (asymmetric energy transfer based on adiabatic and non-adiabatic processes), and we’ll make use of them again when we wish to confirm that our experimental results behave as expected.

Now that we’ve outlined the basic mathematical features and phenomena associated with EPs, it is worth reviewing the evolution of this relatively young field, including both the theoretical understanding and experimental realizations.

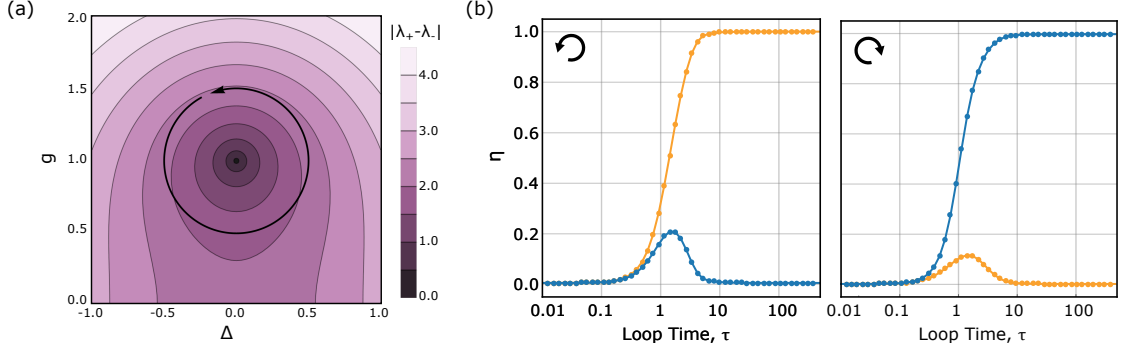


Figure 3.13: Energy transfer dependence on initial condition, perturbation speed, and perturbation direction. (a) Illustration of control loop (counter-clockwise shown), overlaid on a contour plot of the eigenvalue difference, to illustrate the location of the EP. (b) Energy transfer as a function of loop duration, given counter-clockwise (left) and clockwise (right) perturbations, applied to a system initialized in the lower-frequency (orange) or higher-frequency (blue) mode.

3.5 Context and State-of-the-Field

Interest in complex degeneracies and related phenomena grew from several different directions, generally beginning in the late 1980s. This timing was, in part, linked to Michael Berry’s seminal¹³ work [16] on geometric phase in 1984. Within several years of its publication, there was already early work on extending geometric phases to non-Hermitian systems [21]. The possibility of (analytic) eigenmode exchange was noted as early as 1990 [22], and by 1992 it was shown [23] that while adiabaticity cannot be assumed generally, it does hold in the gain manifold (as we saw previously).

Also in the early 1990s, a number of papers began to combine Berry’s concept of cyclic adiabatic processes with descriptions of dissipative atomic levels in terms of complex energies [24, 25]. The result was the realization that optical perturbations of these complex levels could result in transitions from one level to the other. While some of the current language for describing complex degeneracies had not yet been adopted, these papers represent early examples of exploiting EP monodromy to affect useful population transfer. Control of atomic states continues to be a frequently proposed application of EP phenomena [26–28].

In some ways, this early literature represents a rediscovery by the physics community of certain mathematical ideas associated with complex degeneracies, which are of course

¹³ S. Pancharatnam’s work [17] on geometric phases in optical polarization predates this work by almost 30 years, but Berry’s paper seemed to inspire a more broad response from the community.

much older. A thorough mathematical introduction, including use of the term "exceptional points" is usually traced back to a textbook by Tosio Kato in 1966 [29]. The first physicist to pick up this more generic mathematical formalism for complex degeneracies was Dieter Heiss [30]. Heiss was interested in quantum chaos, and how it arises from systems which are dense with avoided level crossings [31]. These crossings became branch point singularities if certain parameters were continued into the complex plane, and Heiss recognized these as exceptional points. Heiss would expand this understanding of the connection between exceptional points, chaos, and phase transitions throughout the 1990s [32]. Heiss also made significant contributions in a pair of papers [33, 34] in which he highlighted several mathematical aspects of EPs, including three which he identified as key characteristics: (1) exchange of eigenmodes upon encircling in the parameter space; (2) a sign change for one of the eigenmodes upon encirclement (implying that it takes four complete loops to return to an initial state); (3) crossing or anti-crossing behavior depending on which 'side' of the EP one considers. These characteristics would become the evaluation criteria for the first experimental demonstration of an EP [35], published in 2001. It is worth spending a moment to review this experiment by Dembowski *et. al.*, in order to understand some of its limitations (which have also been limitations in most of the EP experiments since then).

This experiment [35] focused on creating an EP between two modes of a 3D microwave cavity (Fig. 3.14(a)). A thin copper cylinder (radius 10 cm, height 5 mm) was divided into two semicircular cavities by a thin wall. Each side supported a cavity mode, and these modes could be coupled by opening up a slit in the dividing wall. Changing the width of the slit directly tuned the coupling. A piece of Teflon located in one semicircle allowed for that mode frequency to be independently perturbed. These two parameters (Teflon location, δ , and slit width, s) would constitute the control parameters with which they could perturb their coupled, damped resonators. To demonstrate the modes' eigenvalue topology, they varied δ and s , measuring the complex eigenvalues (mode frequencies and linewidths) at each point. They were able to do this while sweeping δ at two different values of s (one below the EP and one above the EP) and see the expected crossing/anti-crossing behavior in terms of the frequencies and linewidths (Fig. 3.14(b)).

To see the eigenmode switching behavior, the group mapped out the mode shapes at

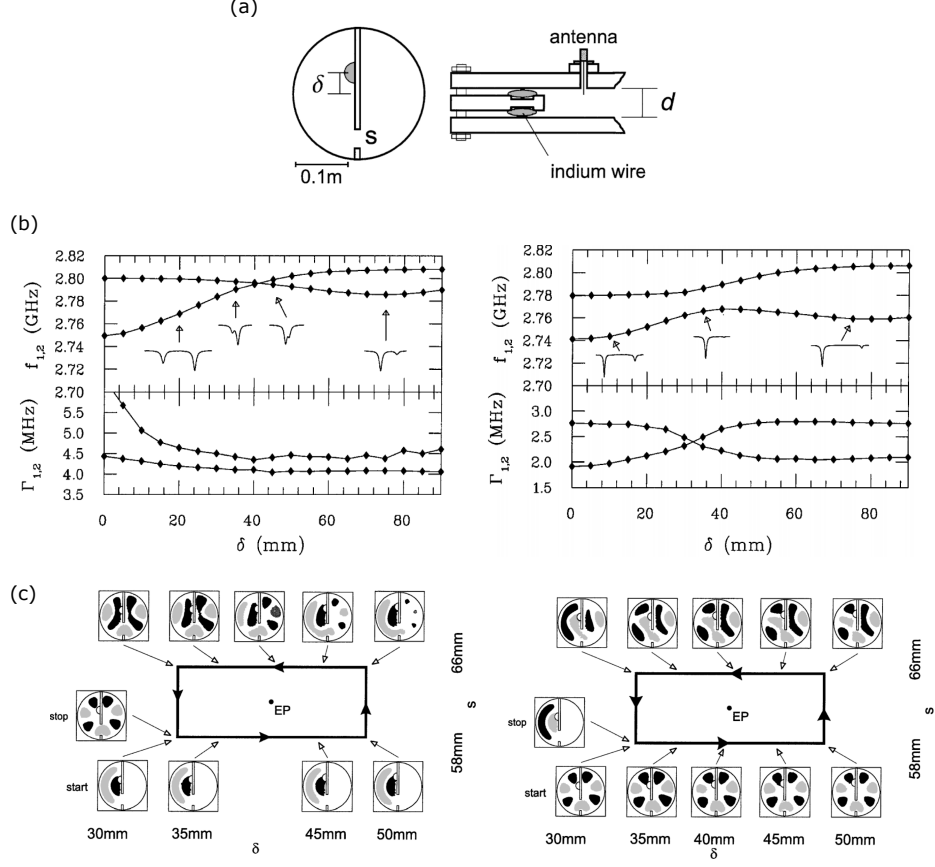


Figure 3.14: Summary of results from [35]. (a) Schematic of the microwave cavity. (b) Eigenvalue measurements (as a function of δ) at two different values of s , showing frequency/linewidth crossing/anti-crossing behavior as expected. (c) Measured field distributions of the two modes (left/right) as the control parameters are varied around the EP. Coloring of nodes/antinodes is chosen arbitrarily at the beginning, then smoothly continued between measurements. Note that the initial mode on the left (right) is the same as the final mode on the right (left) (except for a sign change in one case, as expected)

different values of δ and s , varied along a closed trajectory enclosing the EP. At each point, two different field patterns were measured, and the data shows that along the loop, each mode pattern smoothly transitions to the other (Fig. 3.14(c)).

The last criteria they sought to confirm was the sign change of one of the modes upon encirclement. Observing this directly was not possible, since their measurement could only show them squared field distributions. Instead, they arbitrarily attributed signs to the field distributions at the beginning, then inferred the sign distribution at each step along the loop, based on smooth evolution from one step to the next. Doing so, they were able to confirm that, for example, $\psi_1 \rightarrow \psi_2$ while $\psi_2 \rightarrow -\psi_1$ under application of a counter-clockwise loop (as seen in Fig. 3.14(c)).

This landmark paper was the first demonstration of the topology and chirality expected from exceptional points. However, some of its limitations were immediately apparent in the preceding discussion. For instance, the resonances were indeed tunable, but via a slow process (physically modifying the cavity). This both limited the feasibility of densely mapping the eigenvalues, and prevented any sort of smooth adiabatic manipulation of the parameters¹⁴. These problems would actually be characteristic of some of the other EP experiments conducted in the years after the Dembowski paper. As we'll see in the next chapter, these are precisely the sort of limitations that our optomechanical system is well-poised to avoid.

Since the Dembowski paper, a number of other groups have demonstrated similar measurements of eigenvalues and eigenmode trajectories (again, inferring smooth evolution of the modes as the control parameters are varied). While the conceptual content of these works is not significantly different, it is interesting to note the variety of systems in which EPs have been experimentally explored: photonic crystals [36], chaotic optical microcavities [37], exciton-polariton billiard [38], and even simple analog electrical circuits [39].

Another major area in which there has been significant progress, both experimentally and theoretically, is in the overlap between EP physics and \mathcal{PT} -symmetric systems. Put quite briefly, \mathcal{PT} -symmetric quantum mechanics was proposed as an extension of conventional quantum mechanics [40], in which the requirement that Hamiltonians be hermitian is replaced by the requirement that they be symmetric with respect to an inversion of both spatial and time coordinates. This condition is still sufficient to generate real, positive energies and unitarity, as one would desire from a quantum theory. One realization which has led to a surge in theoretical and experimental work recently is the idea that one can map \mathcal{PT} -symmetric quantum mechanics onto optical systems [41]. The complex potentials one would encounter in \mathcal{PT} -symmetric quantum mechanics are mapped onto optical systems with complex indices of refraction (i.e. gain and loss). As it turns out, the points at which

14. Note that there is a delicate balance necessary to study adiabaticity in open systems. On one hand, perturbations must be slow enough to satisfy some adiabatic-like condition. But at the same time, if the system is subject to loss, then there will be some overall population decay, exponential in the perturbation time. Thus, an important factor in these measurements is the dynamic range of one's readout (i.e. the ability to measure the system before and after it has been subjected to some large decay or amplification).

\mathcal{PT} -symmetry is broken (and real eigenvalues become complex) is precisely an exceptional point. (This is perhaps not entirely surprising, since all of this might be cast under the same umbrella of complex degeneracies). When viewed from the perspective of coupled optical waveguides or resonators, these points where \mathcal{PT} -symmetry is broken (i.e. EPs) lead to a variety of interesting optical phenomena which have been explored in recent years [42, 43].

Finally, the other major conceptual advance for EPs in the last decade has been in understanding adiabaticity (and the lack thereof) in open systems. In 2011, Uzdin *et. al.* provided the first clear description of the asymmetry between propagation along the gain mode and propagation along the loss mode, describing the appearance of non-adiabatic behavior [18]. Around the same time, a paper by Berry and Uzdin [19] considered an analytically solvable model, and made connections between the non-adiabatic phenomena and Stokes theory of asymptotics. These papers represent some of the earliest work which fully addressed the asymmetric behavior we have focused on in this chapter.

An interesting extension of these works was presented by Milburn *et. al.* in 2015 [20]. This work sought to provide a more intuitive mathematical framework for understanding non-adiabatic behavior. The authors describe the behavior in terms of attraction to one of two fixed points corresponding to adiabatic and non-adiabatic evolution. Furthermore, they make connections between temporary evolution along the loss mode and the idea of stability loss delay (a concept from the study of dynamical bifurcations). This paper provides a particularly clear summary of the dynamical behavior one expects near an EP, and was an invaluable resource in motivating our initial foray into this field.

This brief overview of exceptional point literature provides a sense of the state-of-the-art at the time of the experiments presented in this thesis. Broadly speaking, the mathematical properties of complex degeneracies had been understood for several decades, and static measurements of eigenspectra near EPs had been achieved in a number of settings. But the adiabatic phenomena (which were recently becoming much more well-understood) had yet to be demonstrated, since no experiments possessed sufficient real-time tunability. In the next chapter, we'll see how our membrane optomechanical system is remarkably well-suited to demonstrate this dynamical behavior. We will be able to not only demonstrate the static spectrum of an EP in great detail, but actually populate the system and measure

its evolution under application of a time-dependent perturbation.

Chapter 4

Optomechanical Exceptional Points

It was suggested in Ch. 1 that optomechanical systems are well-suited to study the physics of exceptional points, particularly ideas related to adiabatic energy transport. To see this, we will first extend the optomechanical model from Ch. 1 to include multiple mechanical modes, and discuss the resulting coupling between the modes. This coupling will be shown to predict the existence of an EP in our system, with convenient control provided by the optical drive parameters. We then experimentally demonstrate such an EP in our system, including full spectroscopy around the EP, showing the appropriate topology of the eigenvalue manifolds.

Having demonstrated the existence of the optomechanical EP, we then investigate adiabatic energy transfer via closed loops in the control space. I'll first review the measurement protocol for this sort of experiment, then present summary results which illustrate how the energy transfer depends on the properties of the control trajectory. Specifically, we'll see that the system behavior has the correct dependence on the loop geometry, as well as the speed with which we vary the Hamiltonian. Moreover, the energy transfer exhibits an asymmetric dependence on initial condition and loop orientation, consistent with the predictions of Ch. 3.




Figure 4.1: Nearly-degenerate $\{1,3\}/\{3,1\}$ mode pair. (a) Measured optical spectra (proportional to mechanical displacement, plus a noise background). The observed motion (two Lorentzian peaks near $\omega_{1,3}$ and $\omega_{3,1}$) illustrates the modes' mechanical response to a white thermal force noise. (b) 'Ringdown' measurements (red) and fits (black) for measuring the mode decay rates $\gamma_{1,3}$ and $\gamma_{3,1}$. Each mode is mechanically driven, then the drive is turned off and the motion is recorded as it decays (at the mode's intrinsic dissipation rate). Note that the peaks in (a) include a small amount of dynamical backaction (optical spring/damping) from the measurement beam. Thus, the peaks are slightly shifted from $\omega_{1,3}$ and $\omega_{3,1}$, as well as damped. The ringdowns in (b) are measured using an optical wavelength for which the cavity is very low-finesse, so there should be negligible backaction in those measurements.

4.1 Optomechanics with Multiple Mechanical Modes

4.1.1 Motivation for coupled-mode model

Mechanical systems typically support many mechanical modes capable of interacting with the optical field, but as long as their frequencies are well-separated¹, their optomechanical interactions will be independent, and they can be analyzed separately. However, this condition is not always satisfied; in particular, our nearly-square membrane contains nearly-degenerate mode pairs, as discussed in Ch. 2. For instance, in the membrane used in this work, we find that the $\{1,3\}$ and $\{3,1\}$ modes have frequencies $\omega_{1,3}/(2\pi)=788.024$ kHz and $\omega_{3,1}/(2\pi)=788.487$ kHz. The thermal motion of these modes is shown in Fig. 4.1, along with ringdown measurements indicating mechanical damping rates of $\gamma_{1,3}/(2\pi)=0.6$ Hz and $\gamma_{3,1}/(2\pi)=1.4$ Hz. Since the frequency splitting of these modes is only ~ 500 Hz

1. Here, well-separated means with respect to the typical optomechanical coupling strength

(which is smaller than the achievable optomechanical coupling rates), it will be necessary to model these modes in a way that accounts for their interaction and hybridization via the optomechanical coupling. For concreteness in the following discussion, we will focus on this particular pair of modes, since they will eventually be the modes used for studying exceptional points.

4.1.2 Deriving mode coupling

The simple optomechanical model of Ch. 1 is easily extended to include two mechanical modes. We'll pick up the previous derivation at Eq. 1.31, describing the equations of motion for the optical fluctuations d and the complex mechanical amplitudes c . Here, we'll consider two mechanical modes $c_{1,2}$, with frequencies $\omega_{1,2}$, damping rates $\gamma_{1,2}$, and coupling rates $\alpha_{1,2}$. Note that the coupling rates for each mechanical mode in our system can be different, depending on the position of the optical beam relative to their mode profile.

$$\dot{d} = -\left(\frac{\kappa}{2} - i\Delta\right)d - i\alpha_1(c_1 + c_1^*) - i\alpha_2(c_2 + c_2^*) \quad (4.1)$$

$$\dot{c}_1 = -\left(\frac{\gamma_1}{2} + i\omega_1\right)c_1 - i(\alpha_1^*d + \alpha_1d^*) \quad (4.2)$$

$$\dot{c}_2 = -\left(\frac{\gamma_2}{2} + i\omega_2\right)c_2 - i(\alpha_2^*d + \alpha_2d^*) \quad (4.3)$$

As we did for a single mode, we will move to the Fourier domain, substitute $d[\omega]$ into the mechanical equations, and eliminate counter-rotating $c_i^*[\omega]$ terms, yielding

$$\left(\frac{\gamma_1}{2} - i(\omega - \omega_1)\right)c_1[\omega] = (\chi_c^*[-\omega] - \chi_c[\omega])\left(|\alpha_1|^2c_1[\omega] + \alpha_1^*\alpha_2c_2[\omega]\right) \quad (4.4)$$

$$\left(\frac{\gamma_2}{2} - i(\omega - \omega_2)\right)c_2[\omega] = (\chi_c^*[-\omega] - \chi_c[\omega])\left(|\alpha_2|^2c_2[\omega] + \alpha_1^*\alpha_2c_1[\omega]\right) \quad (4.5)$$

For a single mechanical mode, we defined the optomechanical self-energy

$$\Sigma[\omega] = i|\alpha|^2(\chi_c^*[-\omega] - \chi_c[\omega]) \quad (4.6)$$

Here, we will extend this concept to a self-energy matrix:

$$\mathbf{\Sigma} = \begin{pmatrix} i|\alpha_1\alpha_1| & i|\alpha_1\alpha_2| \\ i|\alpha_2\alpha_1| & i|\alpha_2\alpha_2| \end{pmatrix} (\chi_c^*[-\omega] - \chi_c[\omega]) \quad (4.7)$$

Writing our mechanical modes as a vector $\bar{c}[\omega] = \begin{pmatrix} c_1[\omega] \\ c_2[\omega] \end{pmatrix}$, we can write the following matrix equation:

$$-i\omega\bar{c}[\omega] = - \begin{pmatrix} \frac{\gamma_1}{2} + i\omega_1 & 0 \\ 0 & \frac{\gamma_2}{2} + i\omega_2 \end{pmatrix} \bar{c}[\omega] - i\mathbf{\Sigma}[\omega]\bar{c}[\omega] \quad (4.8)$$

Just as we did before, we note that $\mathbf{\Sigma}$ varies with detuning on the scale of κ , while the mechanical modes are only susceptible to drives within their linewidth, which is significantly smaller than κ , by assumption. Therefore, it is sufficient to consider $\mathbf{\Sigma}[\omega] \approx \mathbf{\Sigma}[\omega_1] \approx \mathbf{\Sigma}[\omega_2] \equiv \mathbf{\Sigma}$. (Note that here we have made use of our assumption that the mechanical modes are nearly-degenerate: $\omega_1 - \omega_2 \ll \kappa$). Now that $\mathbf{\Sigma}$ is not a function of ω , we can easily move back to the time domain to find:

$$i\dot{\bar{c}} = \mathbf{H}\bar{c} \quad (4.9)$$

where we define

$$\mathbf{H} = \begin{pmatrix} \omega_1 - i\frac{\gamma_1}{2} & 0 \\ 0 & \omega_2 - i\frac{\gamma_2}{2} \end{pmatrix} + \mathbf{\Sigma} \quad (4.10)$$

Again, it should be emphasized that the entries of $\mathbf{\Sigma}$ are complex, and tunable by the optical drive power P and detuning Δ . It is also worth noting that the entries of $\mathbf{\Sigma}$ have the same complex phase, and differ only in amplitude by their optomechanical coupling strength: $\Sigma_{jk} = i|\bar{a}|^2 g_j g_k (\chi_c^*[-\omega] - \chi_c[\omega])$

We have here an expression which entirely describes our multimode optomechanical system. The two mechanical modes are described by an effective Hamiltonian \mathbf{H} , which includes both their intrinsic resonance parameters as well as an optomechanical modification. That optomechanical interaction consists of diagonal terms (the usual optical

spring/damping for each mode), but also off-diagonal terms which couple the two modes. The result will be a hybridization of the original mechanical modes into new normal modes, which we can study via an eigenvalue analysis of the matrix \mathbf{H} .




Figure 4.2: Optomechanical eigenvalue sheets. Calculated real/imaginary eigenvalues (frequencies/linewidths) of the effective mechanical Hamiltonian. Note that the frequencies and linewidths are viewed from different perspectives, to ensure visibility of the EP. The system parameters are taken from fits described later in the chapter.

4.1.3 Mechanical Eigenmodes: Optical Spring, Damping, and Coupling

Having constructed our effective Hamiltonian matrix \mathbf{H} , and identified P and Δ as our control parameters, we can now simply calculate the spectrum of this matrix as a function of these coordinates. As we did in Ch. 3, we can separately visualize the real and imaginary parts of the spectrum, which correspond to the mechanical frequencies and linewidths, respectively². These eigenvalues are plotted in Fig. 4.2, and we discover one important fact immediately: while there are some distortions, the eigenvalue surfaces in the optomechanical model are qualitatively similar (topologically equivalent, in fact) to the surfaces from the toy model of Ch. 3. For a particular set of control parameters $\{P_{\text{EP}}, \Delta_{\text{EP}}\}$, both the mechanical frequencies and the mechanical linewidths coalesce, and in the immediate vicinity of this degeneracy, both surfaces exhibit the appropriate topology. From this, we can conclude that our optomechanical system indeed possesses an accessible exceptional point³.

2. Note: $\omega_{\pm} = \text{Re}(\lambda_{\pm})$, $\gamma_{\pm} = -2\text{Im}(\lambda_{\pm})$, to accommodate the usual definition of damping rate

3. In fact, it includes two (as it must, since EPs always appear in pairs). If we were to examine the eigenvalues near $\Delta = +\omega_m$ (instead of $\Delta = -\omega_m$), we would find the other EP)

Additional intuition can be had by considering these eigenvalues in the complex plane (as we did for a single optomechanical interaction in Fig. 1.5 and for the eigenvalues of our toy model in Fig. 3.5). Fig. 4.3 shows the same eigenvalues as in Fig. 4.2, but now plotted in this way. The eigenvalues are plotted along contours of constant P , for several powers below and above P_{EP} . Arrows indicate the trajectories of the eigenvalues as Δ is varied from $\Delta=-1100\text{kHz}$ to $\Delta=0$.




Figure 4.3: Calculated optomechanical eigenvalues near an EP. Complex eigenvalues ($\omega_m + i\gamma_m$) plotted along contours of constant P . For a given power, both eigenvalues are plotted in the same color. Arrows indicate the trajectory of the eigenvalues as Δ is varied from $\Delta=-1100\text{kHz}$ to $\Delta=0$. The EP is marked with a black cross. The unperturbed eigenvalues (i.e. when $P = 0 \mu\text{W}$) are shown as black circles. The system parameters are taken from a fit described later in the chapter.

If we first consider the low-power trajectories, we find that these trajectories are qualitatively similar to the single-mode trajectories from Fig. 1.5(a). This is consistent with the idea that the mechanical modes only hybridize if the coupling between them is sufficiently large compared to their separation – for low powers, the modes are essentially non-interacting, and the simple optomechanical model of Ch. 1 (applied separately to each mode) is sufficient to predict their eigenvalue contours.

For high-power trajectories, the behavior is notably different: We see that beyond a certain power, the individual 'loops' merge, resulting in trajectories that begin at one frequency and end at the other. This is already indicative of the eigenvalue-swapping discussed

in Ch. 3. Fig. 4.4 shows a restricted view of these contours near the EP, displaying the same structure we saw in Fig. 3.5.




Figure 4.4: Calculated optomechanical eigenvalues near an EP: Zoom-In. Note the similarity to the eigenvalue contours from the toy model (Fig. 3.5).

4.2 Experimental Measurement of Eigenvalues

Having established that our multimode optomechanical system should possess an EP, we'd now like to see this experimentally. As outlined in Ch. 2, we will make use of our isolated measurement/control capabilities, using the control laser as the source of optomechanical backaction (Σ), and using the measurement laser simply to monitor the eigenmodes and extract their frequencies/linewidths.

To trace out the eigenvalue surfaces, we vary the control beam power and detuning over a grid spanning the range in Fig. 4.2. At each point, we perform a driven response measurement of the mechanical modes near 788 kHz. The mechanical drive is created optically, by applying amplitude modulation (at ω_{drive}) to the control laser. The beat note between these sidebands and the control laser produces a fluctuating radiation pressure force on the membrane at ω_{drive} . This optical amplitude modulation is applied by modulating the power of the RF drive applied to the control beam AOM. We found this optical drive to provide a significantly cleaner result than driving the membrane via the piezo⁴.

4. The piezo seemed to excite other mechanical resonances somewhere in the membrane support structure, which introduced large noise backgrounds underneath the mechanical modes. Optical excitation seems to provide a much more targeted method of driving only the membrane.

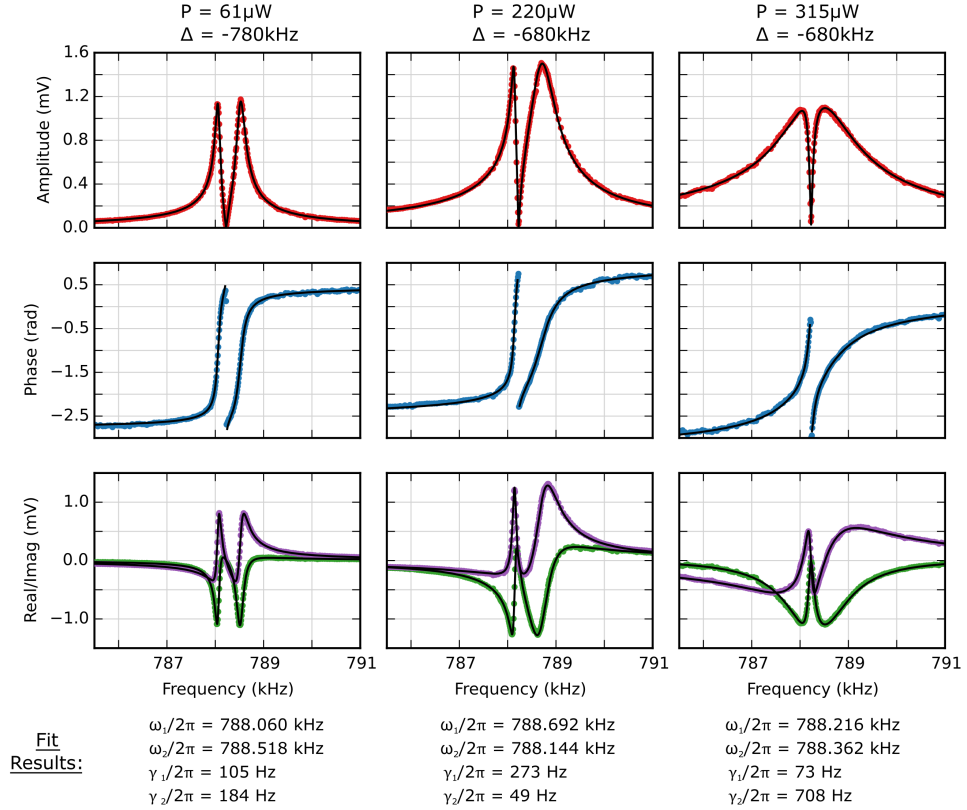


Figure 4.5: Response measurements for extracting mechanical eigenvalues. The lock-in amplifier records the complex response to the drive at each frequency, and here we plot it in several ways (amplitude, phase, in-phase (real) component, out-of-phase (imaginary) component). The real/imaginary components are fit to extract the eigenvalues $\omega_{\pm} + i\gamma_{\pm}$, as described in the main text.

The motion of the membrane is recorded by the measurement beam, and we look specifically at the response to our drive, by applying a narrow-bandwidth ($< 1 \text{ Hz}$) filter to the response, around ω_{drive} . This optically-driven, lock-in measurement of the mechanical response is similar to the technique used in OMIT (Optomechanically Induced Transparency) measurements, except that our drive is being applied to one laser, while the response is measured with another⁵. The lock-in amplifier provides a record of the magnitude and phase of the response, as a function of ω_{drive} ⁶. Fig. 4.5 shows several examples of such measurements, for different values of the control parameters. To extract the mechanical

5. An ambitious author might even describe this as frequency conversion: we inject a signal on the control beam at 280 THz, and use the mechanics to transfer this signal to our measurement beam at 280.008 THz

6. Note that the drive signal itself is coming from the HF2, so we can make a phase-coherent measurement of the response. This drive signal is applied to the modulation input port of the RF signal generator driving the control beam AOM

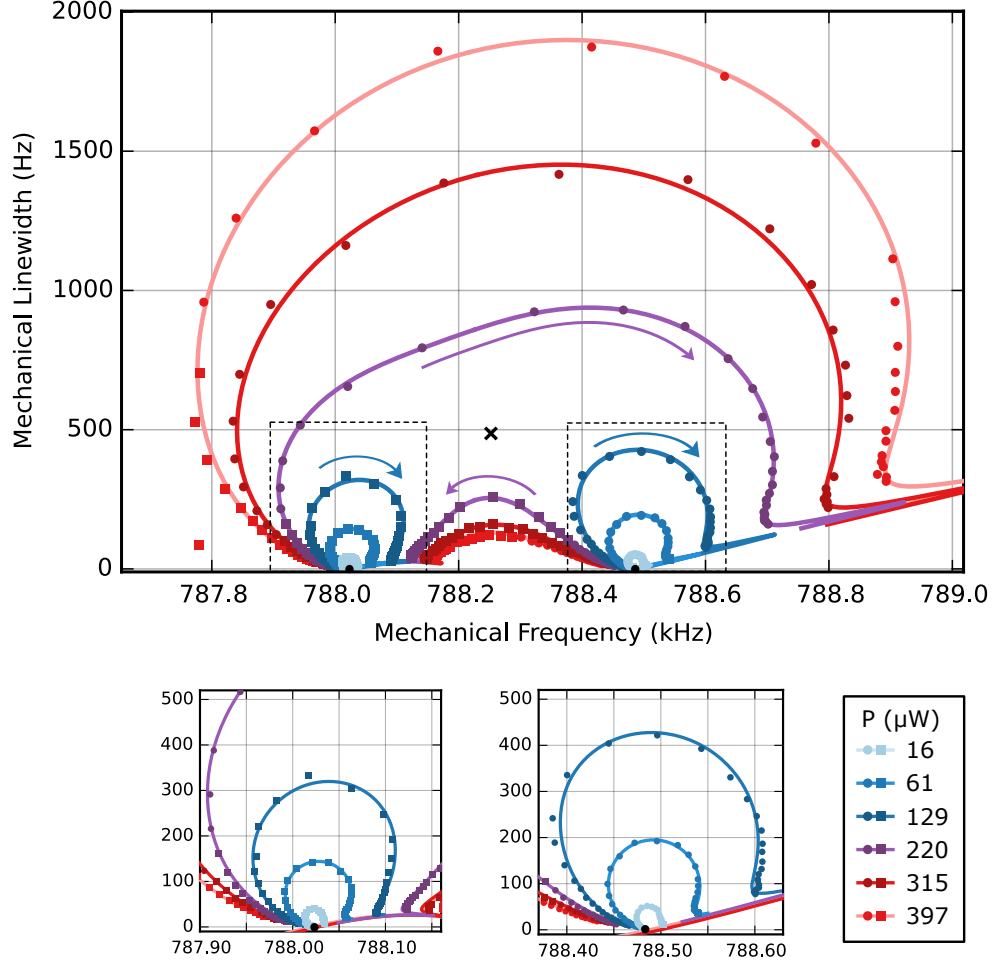


Figure 4.6: Measured mechanical eigenvalues and a least-squares fit to theory. For each power, one set of eigenvalues is plotted in squares, the other in circles. The arrows indicate the direction of eigenvalue evolution as the detuning is varied from negative to positive. The EP is marked with a black cross. The unperturbed eigenvalues are marked with black circles. The solid lines represent a global fit to the data, as described in the main text.

eigenvalues, we fit this data to a generic expression describing two complex resonances:

$$b + \frac{A_+}{\frac{\gamma_+}{2} + i(\omega - \omega_+)} + \frac{A_-}{\frac{\gamma_-}{2} + i(\omega - \omega_-)} \quad (4.11)$$

where each is described by a frequency (ω_{\pm}), a linewidth (γ_{\pm}), and a (complex) amplitude A_{\pm} . A complex background is also allowed as a fit parameter, to account for some non-mechanical response of the circuit. The results of such fits are included in Fig. 4.5, as solid black lines.

Applying this procedure for various P, Δ yields a set of frequencies and linewidths, which

can then be plotted in either of the perspectives discussed in the previous section. In Fig. 4.6, we see one such dataset, consisting of eigenvalues measured along several contours of fixed P , plotted in the complex eigenvalue plane as a function of Δ . This already is qualitatively consistent with the model predictions, but we can go further and fit this data to our model to extract certain system parameters.

The solid lines in Fig. 4.6 represent a single global fit to all of the displayed data, using the bare mechanical frequencies ($\omega_{1,2}$) and linewidths ($\gamma_{1,2}$), optomechanical coupling rates ($g_{1,2}$), and the cavity linewidth (κ) as fit parameters.

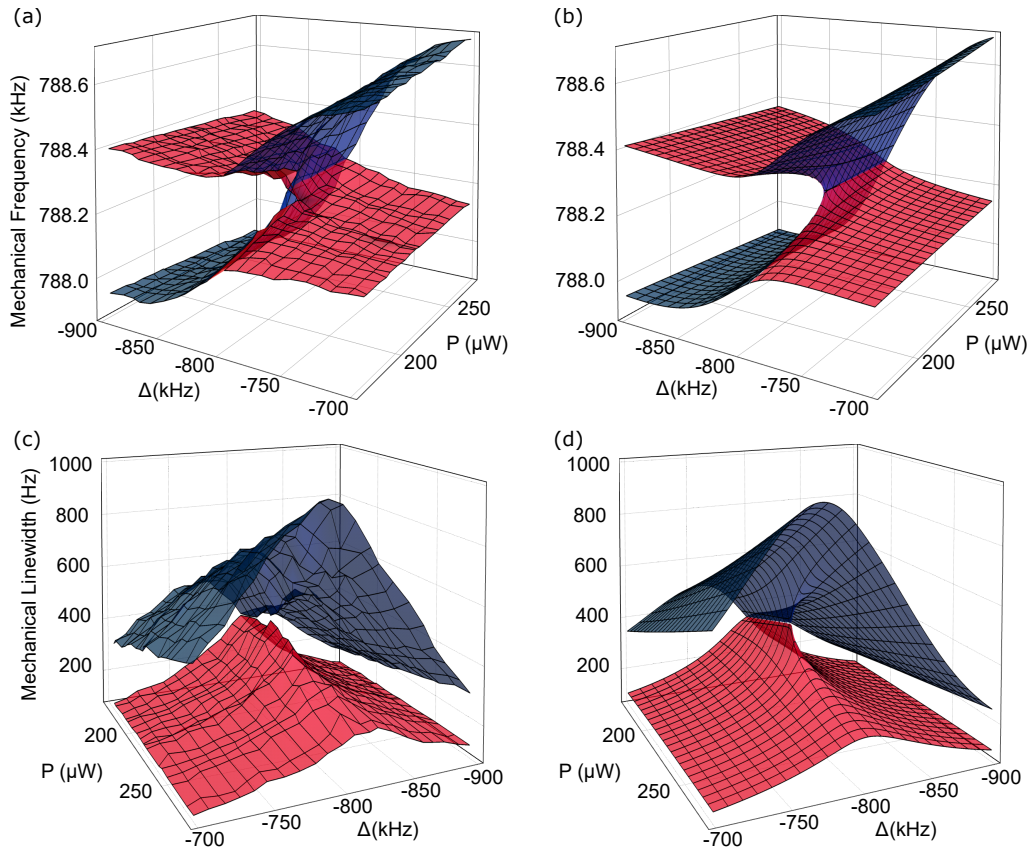


Figure 4.7: Measured (a,c) and theoretical (b,d) eigenvalue sheets. The real/imaginary eigenvalues (frequencies/linewidths) are found by fitting the driven response measurements in Fig. 4.5.

We can also collect a larger, more dense data set near the EP, and plot the frequencies/linewidths against P and Δ , as seen in Fig. 4.7. The intersections of the grid correspond to data points, with surfaces interpolated between the points as a visual aid. For comparison, the predictions of the model are also presented, calculated using the parameters found from the fits in Fig. 4.6.

4.3 Studying Dynamical Behavior: Initialization, Perturbation, and Measurement

As discussed in 3, the eigenvalue topology near an EP leads to interesting phenomena when one considers adiabatic perturbations of the system and closed loops in control space. Having established the existence of an EP in our system, I'll now go beyond static measurements, and demonstrate our ability to study such dynamical effects.

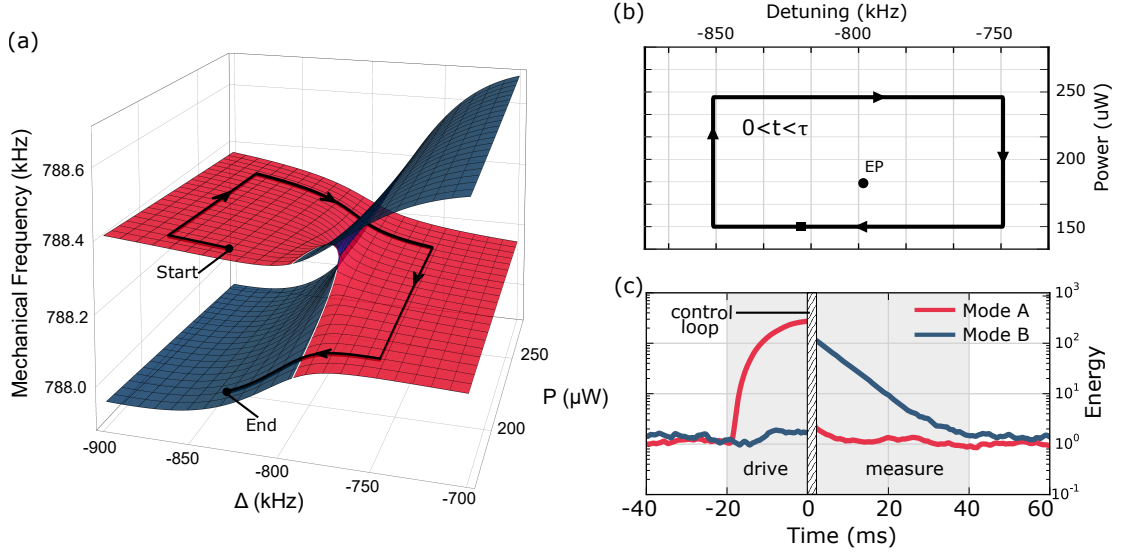


Figure 4.8: Illustration of energy transfer measurement. (a) Eigenvalue trajectory as $\{P, \Delta\}$ are perturbed according to the path in (b), showing a change from one eigenmode to the other. (c) Sample measurement results, in which the higher-frequency mode (red) is driven and a control loop is applied, resulting in a transfer of energy to the lower-frequency (blue) mode. Energy is normalized to the background thermal noise level ($\sim 10^4$ phonons).

The basic idea behind this sort of measurement is qualitatively illustrated in Fig. 4.8. The system is first initialized into one of the eigenmodes, by driving it to a large amplitude (i.e. several orders of magnitude above the usual thermal motion). Then the control parameters (P, Δ) are varied in a closed loop, after which we measure the final distribution of energy between the modes. As discussed previously, one expects to find energy transfer via adiabatic transport for certain conditions (loop time, loop geometry, etc.). Fig. 4.8 shows the result of such a measurement, in which the loop parameters were chosen such that adiabatic energy transport was expected. We will first work through the details of how this sort of measurement is carried out, including both the experimental protocol and

analysis procedure. Then we'll be able to study how this energy transfer depends on loop parameters, and compare with theoretical predictions.

4.3.1 Details of Energy Transfer Measurement

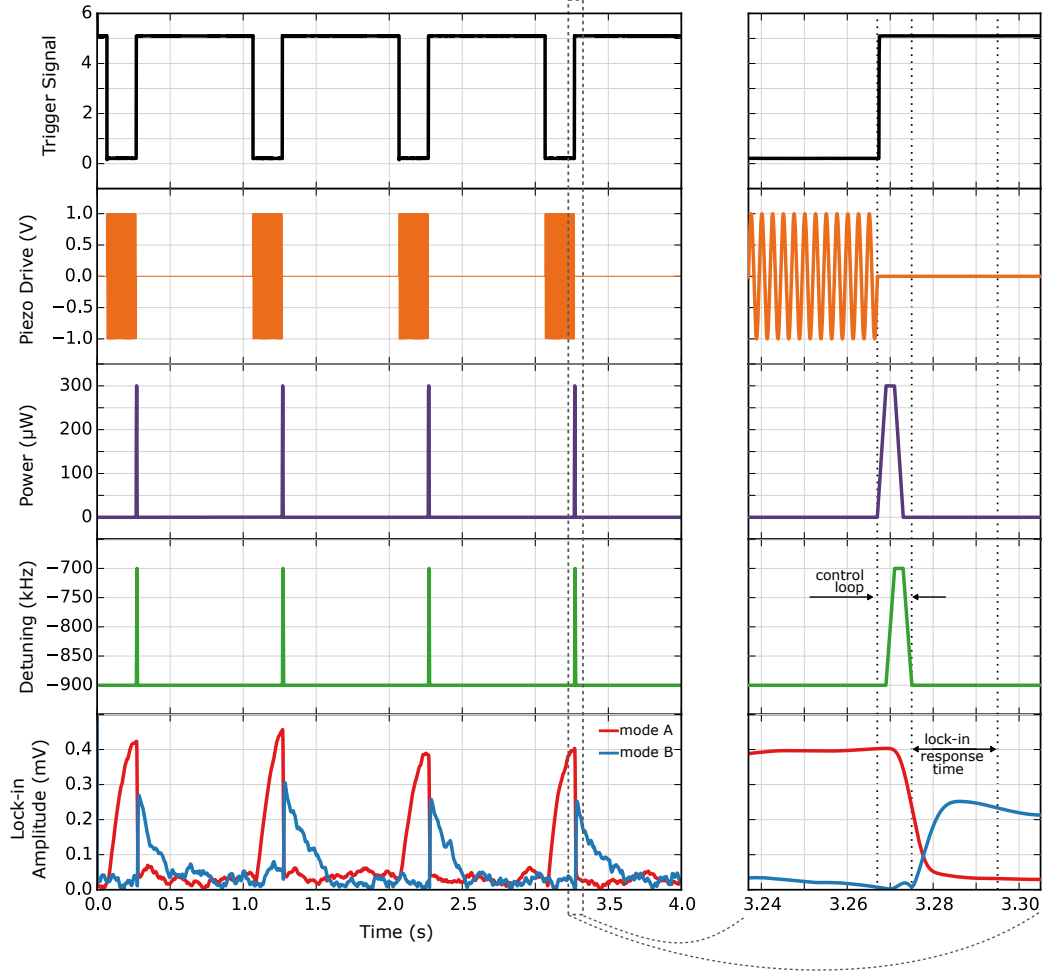


Figure 4.9: Illustration of measurement protocol. From top to bottom, we see the trigger signal (used to synchronize the measurement), the mechanical drive, the power/detuning modulation profiles, and finally the amplitudes of the two eigenmodes throughout the measurement. On the left, several cycles of the measurement are shown. On the right, a zoom-in on one of the control loops is shown. Note that the red/blue curves do not accurately reflect the eigenmode energy during either the control loop or the subsequent lock-in response time. The lock-in amplitudes plotted here are data from an actual measurement, while the other subplots are simulations of the actual signals. Note that the mechanical drive signal shown here is at an artificially low frequency, for illustrative purposes.

A measurement begins with the control beam at a particular $\{P, \Delta\}$, where the mechanical frequencies are $\omega_a(0), \omega_b(0)$. (Note that $\omega_{1,2}$ referred to the bare mechanical frequencies. Here we'll use $\omega_{a,b}(t)$ to refer to the eigenmode frequencies, which depend on time t via

the power and detuning.) We demodulate the heterodyne signal from the measurement beam at these frequencies, using narrow-bandwidth⁷ demodulators which give us separate records of the motion of each mode. The excitation of the state is accomplished by applying a sinusoidal drive (at either $\omega_a(0)$ or $\omega_b(0)$) to the membrane piezo for a fixed amount of time. Then, to create a time-dependent Hamiltonian, we vary P and Δ of the control beam dynamically, during a time τ . This is accomplished by sending particular voltage sequences to the amplitude (AM) and frequency (FM) modulation ports of the RF generator driving the control beam AOM. This will be done in such a way that P and Δ return to their initial values at the end of the sequence. For simplicity, we achieve a loop in $\{P, \Delta\}$ space by a sequence of 4 linear ramps, illustrated in Fig. 4.9. These ramps are provided by Rigol DG1022 signal generators, using their ability to output arbitrary, user-defined functions. One could also easily provide sinusoidal modulation voltages this way, to execute circular loops. Note that during the control loop, the eigenfrequencies of the system are continually changing, while the measurement channels only capture the motion at $\omega_a(0)$ and $\omega_b(0)$. Thus, these signals are not meaningful during the control loop, but since $\omega_{a,b}(0) = \omega_{a,b}(\tau)$, they do provide an accurate record of the energy distribution before and after the loop.

It is important that the drive, control loop, and measurement records are all well-synchronized for this type of measurement sequence. Fig. 4.10 illustrates the circuit used to ensure this synchronization. A clock for the whole process is provided by a rectangle wave generated by RSG1. This clock, in part, provides the control voltage for an RF switch, which gates the sinusoidal mechanical drive provided by the HF2. When the clock signal is low, the drive is on, exciting one of the membrane modes. The low-time for the clock is determined by the time necessary to excite the mechanical mode to the desired amplitude. The clock is also fed into RSG2 and RSG3, as an external trigger. When the clock goes high, it simultaneously cuts the mechanical drive and triggers RSG2 and RSG3 to output their user-defined voltage ramps. These ramps translate into the looping of P and Δ for

7. The size of this bandwidth Δf represents a trade-off between competing factors. The minimum time for the channel to respond to changes in the signal is $\frac{1}{\Delta f}$, so the bandwidth should not be too small. At the same time, if Δf is too large, a greater amount of thermal background will be included, potentially even including motion from the other mode. In practice, we choose $\Delta f=50$ Hz, which results in a response time of 20 ms. This time is accounted for in our later analysis

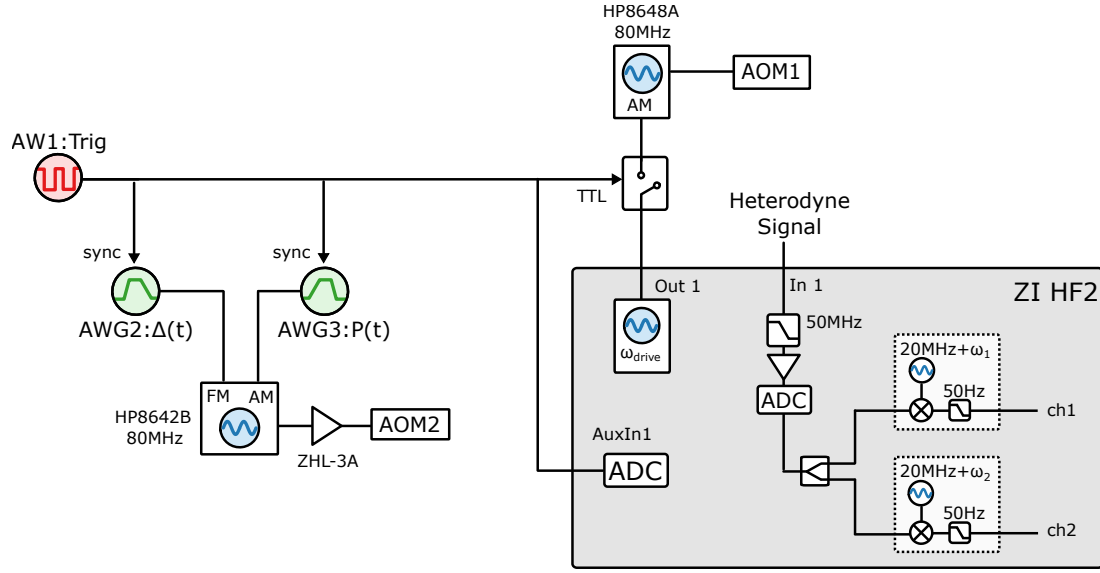


Figure 4.10: Electric circuit for energy transfer measurements. The trigger from AWG1 is used to trigger both the mechanical drive (applied via AOM1) and the power/detuning modulation (applied via AWG2 and AWG3). The HF2 simultaneously provides the mechanical drive signal and demodulates the heterodyne signal in a narrow bandwidth near the eigenmode frequencies

the control beam. The clock should remain high long enough for the control loop as well as whatever time is necessary for the driven motion in the modes to decay, so that the system returns to its equilibrium state, ready for the process to be repeated. Finally, the clock is also collected by the HF2, for synchronization of the data stream in post-processing.

To increase the signal-to-noise of our measurement, this process of initialization, cyclic perturbation, and measurement is repeated many times, as shown in Fig. 4.9, and the amplitudes of the samples are averaged together. The clock trigger signal is used to identify individual measurements and align them, then their amplitudes are averaged together to produce the dark solid lines in Fig. 4.11.

It is also necessary to convert this voltage signal into mechanical amplitude (or energy). In particular, it is important to apply this conversion because the signal from each mode is scaled by its optomechanical coupling rate $g_{1,2}$. This conversion factor should account for all stages of the measurement process: transduction of motion onto optical phase, cavity filtering of this phase modulation, attenuation of the optical signal before reaching the photodiode, measurement of optical phase via heterodyne interferometry, and attenuation by the electrical measurement circuit. All of these factors have been summarized in a previous




Figure 4.11: Illustration of measurement averaging procedure. The faint red/blue lines are the results of individual energy transfer measurements. The solid, dark red/blue lines represent the averaged result.

work [10], which derives the heterodyne spectrum of a mechanical oscillator equilibrated with a thermal bath. With some algebraic manipulation, this result can be interpreted as a scaling between the signal voltage (V_{sig}) and energy. Due to the possibility of different electrical gain ($G_{r,b}$) and cavity filtering, the red and blue sidebands have slightly different scaling:

$$\text{Red Sideband : } E_{\text{mech}} = \hbar\omega_m V_{\text{sig}}^2 / \left(\sigma^2 G_r^2 \beta^2 \frac{P_{\text{LO}}}{\hbar\omega_L} \kappa_{\text{in}} |\alpha_{\text{meas}}|^2 |\chi_c[-\omega_m]|^2 \right) \quad (4.12)$$

$$\text{Blue Sideband : } E_{\text{mech}} = \hbar\omega_m V_{\text{sig}}^2 / \left(\sigma^2 G_b^2 \beta^2 \frac{P_{\text{LO}}}{\hbar\omega_L} \kappa_{\text{in}} |\alpha_{\text{meas}}|^2 |\chi_c[+\omega_m]|^2 \right) \quad (4.13)$$

where V_{sig} is the signal amplitude, P_{LO} is the local oscillator power, ω_L is the laser frequency, α_{meas} is the multi-photon coupling rate of the measurement beam (defined as in Ch. 1).

Applying this factor (and accounting for the differing $g_{1,2}$ between the modes), we now have records of mechanical energy as a function of time for each mode. Note that there is a non-zero background to which both modes decay, coming from the thermal motion of the membrane. In all the data considered going forward, this will not significantly affect our results. In all measurements, we initially excite one mode such that the driven coherent motion far exceeds the thermal motion (by several orders of magnitude). In this way, the




Figure 4.12: Quantitative analysis of energy transfer. Here, $t = 0$ corresponds to the beginning of the control loop. The first dashed line shows the end of the control loop (τ) and the second dashed line indicates the time constant of the measurement channel. Since the measurement channel cannot respond to changes faster than this, we only begin fitting the data at this point. However, we infer the energy at τ by extrapolating the exponential fit backwards, as shown. The black dots indicate the energy values used in calculating propagation matrices and energy transfer efficiencies.

initial state vector can be well-approximated by

$$\begin{pmatrix} c_a(0) \\ c_b(0) \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \text{ or } \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (4.14)$$

The data plotted in Figs. 4.8 and 4.11 depict a measurement in which the loop size, orientation, and initial condition were chosen so as to demonstrate energy transfer (i.e. by encircling the EP and being sufficiently slow). Moving forward, our goal will be to quantitatively analyze this energy transfer as various conditions are changed. Therefore, we now introduce a procedure for extracting quantitative information from this data.

The key information we need is the energy of the eigenmodes after the control loop is complete: $E_{a,b}(\tau)$. Since our measurement record contains some noise, we choose to fit the energy over a wide range of $t > \tau$, and infer $E_{a,b}(\tau)$ from this fit. For $t > \tau$, the energy of each mode is described by a simple exponential relaxation to its thermal equilibrium. The start of the fit window is intentionally offset from τ by 20ms, to allow time for the the 50 Hz measurement bandwidth to respond to the energy transfer during the control loop. For comparison to the theoretical prediction of the state amplitudes, we extrapolate this exponential fit backwards to τ , as seen in Fig. 4.12.

As we proceed, there will be two primary metrics we use to describe the energy dynamics during a control loop. In some situations, we will simply describe the control loop in terms of a propagation matrix, which maps the initial state to the final state:

$$\begin{pmatrix} c_a(\tau) \\ c_b(\tau) \end{pmatrix} = \mathbf{U}^{\text{control}} \begin{pmatrix} c_a(0) \\ c_b(0) \end{pmatrix} \quad (4.15)$$

By performing separate measurements in which we initialize in either mode a or b , then measure the final energy in both modes, we can calculate the amplitudes of this propagation matrix. For example, by first initializing in mode a , we can measure

$$\left| \mathbf{U}_{11}^{\text{control}} \right| = \frac{|c_a(\tau)|}{|c_a(0)|} = \sqrt{\frac{E_a(\tau)}{E_a(0)}} \quad \text{and} \quad \left| \mathbf{U}_{21}^{\text{control}} \right| = \frac{|c_b(\tau)|}{|c_a(0)|} = \sqrt{\frac{E_b(\tau)}{E_a(0)}} \quad (4.16)$$

Then, by initializing in mode b , we can measure

$$\left| \mathbf{U}_{12}^{\text{control}} \right| = \frac{|c_a(\tau)|}{|c_b(0)|} = \sqrt{\frac{E_a(\tau)}{E_b(0)}} \quad \text{and} \quad \left| \mathbf{U}_{22}^{\text{control}} \right| = \frac{|c_b(\tau)|}{|c_b(0)|} = \sqrt{\frac{E_b(\tau)}{E_b(0)}} \quad (4.17)$$

Alternatively, we will sometimes summarize a particular control loop and initial condition more succinctly, via a metric we'll call the 'transfer efficiency' η , defined as the fraction of energy at τ which is not in the initial mode. (Note that the utility of this metric is dependent on the fact that we always initialize entirely in one mode or the other). For example, if we initialize in mode a , then $\eta_{a \rightarrow b} = \frac{E_b(\tau)}{E_a(\tau) + E_b(\tau)}$, and if we initialize in mode b , then $\eta_{b \rightarrow a} = \frac{E_a(\tau)}{E_a(\tau) + E_b(\tau)}$. Dividing by the total energy at τ allows us to normalize away the total energy loss during the control loop.

4.4 Analysis: Varying Loop Speed and Initial Condition

Having outlined the basic measurement and analysis procedure, we can now consider varying the parameters of the control loop to see if the system behaves as predicted in Ch. 3. We start by considering how the system responds to perturbations of different speeds, and in

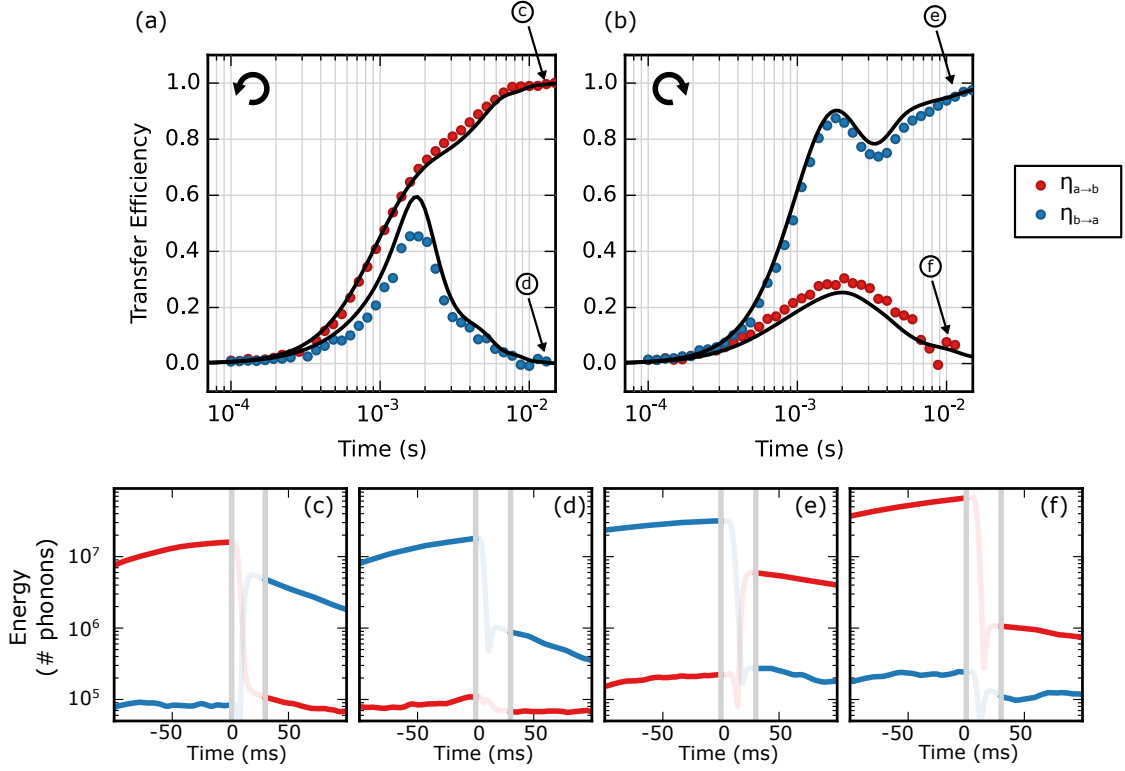


Figure 4.13: Dependence of energy transfer on adiabaticity, initial condition, and loop direction. (a) and (b) show how the transfer efficiency depends on the loop time for counter-clockwise and clockwise loops, respectively. In each, the color of the points (red/blue) indicates which mode was driven initially. (Red corresponds to $\omega_a(0)$, blue corresponds to $\omega_b(0)$) The four insets depict time samples from each data set in the long-time limit, plotted just before and after the control loop. The gray lines indicate the window in which we cannot accurately estimate the eigenmode amplitudes.

turn, how this behavior varies based on initial condition and loop direction.

In all of the following, we work with a 'large' control loop, enclosing the EP. (The sense in which we mean 'large' will be made clear in Sec. 4.5, in which we analyze the dependence on loop geometry.) The loop itself is given by a rectangle in P, Δ space, defined by the points $P_{\min} = 2 \mu\text{W}$, $P_{\max} = 622 \mu\text{W}$, $\delta_{\min} = -1890 \text{ kHz}$, $\Delta_{\max} = -290 \text{ kHz}$. As discussed in Ch. 3, the topology of the EP is only manifest in energy transfer if the system adiabatically follows the instantaneous eigenvalue manifold. This requires an adiabatic perturbation, and hence we expect to see this energy transfer emerge only as the loops become sufficiently slow. Moreover, we expect that the relative loss/gain in the system should mean that this adiabatic transfer between eigenmodes only occurs for certain loop orientations and initial conditions. Specifically, given loops of a particular orientation, we expect that the transfer

efficiency from one eigenmode should approach 0 in the adiabatic limit, while the transfer efficiency from the other eigenmode should approach 1. Alternatively, given loops beginning in a particular eigenmode, we expect the transfer efficiency in one direction to approach 0 in the adiabatic limit, while the transfer efficiency in the other direction to approach 1.

Fig. 4.13 shows the transfer efficiency measured as a function of loop time for each of the four combinations of initial eigenmode and control loop orientation. Qualitatively, we can immediately see that no energy transfer occurs for ‘sudden’ perturbations ($\tau \sim 100 \mu\text{s}$), while in the adiabatic limit ($\tau \gtrsim 10 \text{ ms}$), we see the expected adiabatic behavior. Moreover, the presence of adiabatic energy transfer depends on both the initial condition and the loop direction, as predicted⁸.

The black lines in Fig. 4.13 are the result of numerically simulating the system using Mathematica. This simply involves numerically integrating Eq. 4.9 between $t = 0$ and $t = \tau$ given a particular initial condition, and then calculating the distribution of energies at $t = \tau$. We see that for all four datasets, the numerical results closely match the experimental results.

For these same data, we can take a step backwards in abstraction, and plot the propagator matrix element amplitudes directly, rather than combining them into the transfer efficiency metric. This view of the data is presented in Fig. 4.14. This is useful for highlighting a number of features of the data. First, all of our discussion of the system’s asymmetric dependence on initial condition and loop direction can be summarized by noting that the propagator matrix $\mathbf{U}^{\text{control}}$ for a particular control loop should be (and is measured to be) non-reciprocal – that is, $\mathbf{U}_{12}^{\text{control}} \neq \mathbf{U}_{21}^{\text{control}}$. (While we technically only present the absolute value of these matrix elements, that is sufficient to confirm the inequality.) Moreover, we expect the matrix elements of a clockwise loop and a counterclockwise loop to have a particular relationship, given that one control loop is simply a time-reversal of the other. Specifically, it can be shown that the propagators should satisfy $U_{\circ} = U_{\circ}^T$. Again, we are only measuring the amplitudes of these matrix elements, but the data is consistent with

8. As mentioned in Ch. 3, there are parameter regimes in which this prediction changes to reflect the very-long-term dominance of the gain mode. This does not fundamentally alter the asymmetries or non-reciprocity discussed in this data. These very-long-time corrections will be discussed in Appendix. A.




Figure 4.14: Amplitude of propagator matrix elements. Results of (counter-)clockwise control loops are shown in circles (squares). For each control loop direction, the inequality of the red/blue markers in the long-time limit indicate the non-reciprocity of the propagator matrix. The equality of green (purple) circles with green (purple) squares, as well as the equality of red (blue) circles with blue (red) squares supports the statement that $U_O = U_O^T$

this relationship.

4.5 Analysis: Dependence on Loop Geometry

As discussed in Ch. 3 energy transfer associated with encirclement of an EP is fundamentally related to the topology of the eigenvalues: Loops that adiabatically encircle the EP will exchange eigenvalues, while small loops which do not encircle the EP will not. To see this, we will consider loops of a fixed time and initial condition, and vary the size of the loops between these extremes. For example, the rectangular loop in $\{P, \Delta\}$ space can have one of its edges gradually extended, as depicted in Fig. 4.15. We'll first consider expanding the maximum detuning of the loop, and measuring the variation in transfer efficiency. For all measurements, we initially excite mode 'a' and execute a counter-clockwise control loop with a duration⁹ of $\tau = 16$ ms. In Fig. 4.15, we see qualitatively the behavior we expect: small loops, not enclosing the EP have negligible energy transfer, while large loops enclosing the

9. Note that the loop time is kept constant as the size is varied, and hence the rate of change of the control parameters is not necessarily constant. In this sense, the loop geometry is not the only factor which varies between data points, but we must make some choices in terms of what we choose to hold constant.




Figure 4.15: Dependence of energy transfer on loop geometry. (a) Varying the size of the loop by changing its maximum detuning extent. (b) Varying the size of the loop by changing its maximum power extent. In all measurements, the loop time is 16ms. The initial/final point of the loop is marked by a black square.

EP have near-unity transfer efficiency. As discussed in Ch. 3, the exact behavior when the loop edge passes near the EP will be complicated by the fact that this fixed τ will no longer satisfy adiabaticity when the eigenvalues become very close. Fig. 4.15 also shows a separate measurement in which the control loops are stretched along the P axis, demonstrating the same behavior. As in Fig. 4.13, the black lines are the result of numerical simulations of the system.

Chapter 5

Virtual Exceptional Points from Non-Degenerate Modes

In the previous chapter, our ability to explore the physics associated with exceptional points was predicated on the existence of a convenient symmetry in our system. Namely, the near-symmetry of our membrane resulted in the existence of nearly-degenerate mechanical modes which we could easily couple through a common optical interaction. However, if the bare frequency splitting (~ 500 Hz for those modes) had been larger, our ability to access the EP and demonstrate topological energy transfer would have been limited for several key reasons. First, a larger bare splitting would require larger optomechanical coupling strength to reach the EP – potentially an infeasibly large coupling. Moreover, this larger coupling strength would translate into a larger mechanical linewidth at the EP. Thus, a loop encircling the EP would traverse regions of significantly larger dissipation, leading to greater total loss during the operation. While not a fundamental problem, this again presents a technical difficulty, given the finite signal-to-noise ratio of our measurement.

Despite the preceding discussion, we'll see in this chapter that topologically-protected energy transfer can in fact be achieved with non-degenerate modes, given a modified optical drive scheme. Specifically, we'll consider driving one cavity mode with two control lasers, with carefully chosen detunings, as illustrated in Fig. 5.1. We'll begin by working through the theoretical basis for this type of drive scheme, and find that it results in a coupling of

the non-degenerate mechanical modes. Moreover, we'll find that in a particular reference frame, these modes can be treated as nearly-degenerate, with a frequency splitting which is now under our control. With this coupling and near-degeneracy, we can repeat the experiments of the previous chapter, and find the same sort of eigenvalue topology and dynamical behavior. Thus, we will demonstrate that the topologically-protected energy transfer of the previous chapter need not be limited to nearly-degenerate modes, and so can be achieved in a more generic multimode system.

5.1 Coupling through a Bichromatic Drive: Motivation




Figure 5.1: Frequency layout of bichromatic coupling scheme. By setting the frequency splitting between the lasers (Δ_{12}) to be nearly equal to the difference between the mechanical mode frequencies ($\omega_a - \omega_b$) the ω_b sideband from laser 1 can be made to overlap with the ω_a sideband from laser 2, leading to coupling as described in the text. Note that all sidebands plotted here are anti-Stokes (‘blue’) sidebands – the ‘red’ Stokes sidebands of each laser are not shown. In principal, there is a similar overlapping of the Stokes sidebands, though the contribution to the coupling is suppressed by the cavity transfer function (since we are working in the resolved-sideband regime). This does mean, however, that the same coupling can be generated by two blue-detuned lasers (or even one red-detuned laser and one blue-detuned laser).

We begin by briefly sketching the idea behind this two-laser drive scheme. Consider two mechanical modes with frequencies ω_a and ω_b , with $\omega_a < \omega_b$. Now consider driving a cavity with one laser detuned to $\Delta_1 \approx -\omega_a$ and another detuned to $\Delta_2 \approx -\omega_b$, as illustrated in Fig. 5.1. For each laser, the figure shows the anti-Stokes sidebands from each mechanical mode, and we see that by choosing $\Delta_1 - \Delta_2 \approx \omega_b - \omega_a$, some of these sidebands will overlap. In particular, near the cavity resonance, we see the ω_a sideband from laser 1 overlapping with the ω_b sideband from laser 2. For intuition as to how this can generate coupling between the modes, consider the beat notes formed when these particular mechanical sidebands mix with the laser drives: The optical drive at Δ_1 will beat with the ω_b sideband at a frequency

near ω_a , and the optical drive at Δ_2 will beat with the ω_a sideband at a frequency near ω_b . Thus, the motion of the $a(b)$ mode becomes a force on the $b(a)$ mode – hence coupling the two modes.

5.2 Coupling through a Bichromatic Drive: Derivation

Now, we'll work through the optomechanical equations of motion for this bichromatic drive scheme, to see how this coupling can be described analytically. As in the simple model above, we consider two non-degenerate mechanical modes with frequencies ω_a , ω_b and damping rates γ_a , γ_b . Without loss of generality, we'll choose $\omega_a < \omega_b$. Here, 'non-degenerate' means that we are assuming $|\omega_a - \omega_b|$ is much larger than any of the optomechanical coupling rates (this condition will be made concrete later in the derivation). These mechanical modes are coupled to a single optical mode with coupling rates g_a and g_b . That optical mode (at frequency ω_c) is driven by two coherent inputs at frequencies ω_1 and ω_2 (where we are assuming, without loss of generality, that $\omega_1 > \omega_2$). The coupled equations governing this system follow in the same manner as Ch. 1 and 4:

$$\dot{a} = -\left(\frac{\kappa}{2} + i\omega_c\right)a - i \sum_{k=a,b} g_k(c_k + c_k^*)a + \sqrt{\kappa_{\text{in}}}a_{\text{in}}(t) \quad (5.1)$$

where, defining detunings $\Delta_i = \omega_i - \omega_c$, we can write the input field as

$$a_{\text{in}}(t) = e^{-i\omega_c t} (a_{\text{in},1}e^{-i\Delta_1 t} + a_{\text{in},2}e^{-i\Delta_2 t}) \quad (5.2)$$

For a single cavity drive, we linearized around a mean field, \bar{a} , which was static in the frame of the laser. This time, we'll work in the frame of the cavity, and linearize around $\bar{a}(t) = e^{-i\omega_c t} (\bar{a}_1 e^{-i\Delta_1 t} + \bar{a}_2 e^{-i\Delta_2 t})$, where $\bar{a}_n = \frac{\sqrt{\kappa_{\text{in}}}a_{\text{in},n}}{\frac{\kappa}{2} - i\Delta_n}$. Let's first explicitly check that this solves the bare cavity equation by substituting $\bar{a}(t)$ for $a(t)$ (with the mechanical interaction ignored). First, the left hand side:

$$\dot{a} = \frac{d}{dt} (e^{-i\omega_c t} (\bar{a}_1 e^{-i\Delta_1 t} + \bar{a}_2 e^{-i\Delta_2 t})) \quad (5.3)$$

$$= e^{-i\omega_c t} (-i\omega_c (\bar{a}_1 e^{-i\Delta_1 t} + \bar{a}_2 e^{-i\Delta_2 t}) - i\Delta_1 \bar{a}_1 e^{-i\Delta_1 t} - i\Delta_2 \bar{a}_2 e^{-i\Delta_2 t}) \quad (5.4)$$

$$= e^{-i\omega_c t} (-i\omega_1 \bar{a}_1 e^{-i\Delta_1 t} - i\omega_2 \bar{a}_2 e^{-i\Delta_2 t}) \quad (5.5)$$

$$= e^{-i\omega_c t} \left(-i\omega_1 \frac{\sqrt{\kappa_{\text{in}}} a_{\text{in},1}}{\frac{\kappa}{2} - i\Delta_1} e^{-i\Delta_1 t} - i\omega_2 \frac{\sqrt{\kappa_{\text{in}}} a_{\text{in},2}}{\frac{\kappa}{2} - i\Delta_2} e^{-i\Delta_2 t} \right) \quad (5.6)$$

$$(5.7)$$

Substituting on the right hand side, we have

$$= -\left(\frac{\kappa}{2} + i\omega_c\right) a + \sqrt{\kappa_{\text{in}}} a_{\text{in}}(t) \quad (5.8)$$

$$= e^{-i\omega_c t} \left(-\left(\frac{\kappa}{2} + i\omega_c\right) \left(\frac{\sqrt{\kappa_{\text{in}}} a_{\text{in},1}}{\frac{\kappa}{2} - i\Delta_1} e^{-i\Delta_1 t} + \frac{\sqrt{\kappa_{\text{in}}} a_{\text{in},2}}{\frac{\kappa}{2} - i\Delta_2} e^{-i\Delta_2 t} \right) \right) \quad (5.9)$$

$$+ e^{-i\omega_c t} \sqrt{\kappa_{\text{in}}} (a_{\text{in},1} e^{-i\Delta_1 t} + a_{\text{in},2} e^{-i\Delta_2 t}) \quad (5.10)$$

$$= e^{-i\omega_c t} \left(\frac{-i\omega_1}{\frac{\kappa}{2} - i\Delta_1} \sqrt{\kappa_{\text{in}}} a_{\text{in},1} e^{-i\Delta_1 t} + \frac{-i\omega_2}{\frac{\kappa}{2} - i\Delta_2} \sqrt{\kappa_{\text{in}}} a_{\text{in},2} e^{-i\Delta_2 t} \right) \quad (5.11)$$

Comparing these expressions, we see that our definition of $\bar{a}(t)$ is indeed a good solution for the bare cavity field. Now, to account for the mechanics, we'll rewrite our cavity field in terms of fluctuations $d(t)$ around this field, as we did before: $a(t) = e^{-i\omega_c t} (\bar{a}(t) + d(t))$. Substituting this into Eq. 5.1, we'll of course find that many of the \bar{a} terms will cancel out, leaving us with the following:

$$\dot{d} = -\frac{\kappa}{2} d - i \sum_{k=a,b} g_k (\bar{a} + d) (c_k + c_k^*) \quad (5.12)$$

Just as in the single-mode case, we will drop the $dc_k^{(*)}$ terms, since d and $c_k^{(*)}$ are both small terms, yielding:

$$\dot{d} = -\frac{\kappa}{2} d - i \sum_{k=a,b} g_k \bar{a} (c_k + c_k^*) \quad (5.13)$$

Next, we substitute our expression for $\bar{a}(t)$ into the mechanical equation, to find:

$$\dot{c}_k = -\left(\frac{\gamma_k}{2} + i\omega_k\right) c_k - g_k \left(|\bar{a}(t)|^2 + \bar{a}(t)d^*(t) + \bar{a}^*(t)d(t) + |d(t)|^2 \right) + \sqrt{\gamma_k}\eta_k \quad (5.14)$$

where, as before, η_k indicates a thermal Langevin force driving mode c_k . The $|\bar{a}(t)|^2$ term can be expanded as:

$$|\bar{a}(t)|^2 = \underbrace{|\bar{a}_1|^2 + |\bar{a}_2|^2}_{\text{Static forces}} + \underbrace{\bar{a}_1\bar{a}_2^*e^{-i(\Delta_1-\Delta_2)t} + \bar{a}_1^*\bar{a}_2e^{-i(\Delta_2-\Delta_1)t}}_{\text{Fluctuating force at } |\Delta_1-\Delta_2|} \quad (5.15)$$

Both of these contributions can be ignored: the first is simply a constant force, resulting in a static displacement (just as we find in the single-mode case), and the second represents a time-dependent force at frequency $|\omega_1 - \omega_2|$. Here, we make another important assumption, which is that $|\Delta_1 - \Delta_2|$ is not resonant with either of the mechanical modes (i.e. $|\Delta_1 - \Delta_2| - \omega_k \gg \tilde{\gamma}_k$, where $\tilde{\gamma}_k$ is the effective linewidth of the mode, including all optical contributions). This means that the beat note between our two optical drives will not mechanically drive either of the modes. In general, we are going to consider the situation where $\Delta_1 - \Delta_2 \approx \omega_a - \omega_b$, so this condition will be satisfied as long $\omega_b \neq 2\omega_a$. (So, for example, the experiments carried out below could not be carried out with the (1,1) and (2,2) modes of the membrane, since their frequencies differ by a factor of 2.) In fact, one should avoid using any pair of symmetric modes: since $\omega_{(j,j)} = j\omega_{(1,1)}$, the frequency difference between any pair of symmetric modes will necessarily match one of the other symmetric mode frequencies. Therefore, even if the $|\Delta_1 - \Delta_2|$ beat note did not drive ω_a or ω_b , it would drive some other mode of the membrane, potentially to an amplitude which would interfere with the measurement.

So we can ignore this $|\bar{a}(t)|^2$ term as well as the $|d(t)|^2$ term, since we are keeping only terms which are linear in small fluctuations. This leaves us with the following system of equations for the optical fluctuations and mechanical modes:

$$\begin{cases} \dot{d} = -\frac{\kappa}{2}d - i \sum_{k=a,b} g_k \bar{a} (c_k + c_k^*) \\ \dot{c}_k = -\left(\frac{\gamma_k}{2} + i\omega_k\right) c_k - ig_k (\bar{a}(t)d^*(t) + \bar{a}^*(t)d(t)) + \sqrt{\gamma_k}\eta_k \end{cases} \quad (5.16)$$

If we now substitute $\bar{a}(t) = \bar{a}_1 e^{-i\Delta_1 t} + \bar{a}_2 e^{-i\Delta_2 t}$, we have:

$$\begin{cases} \dot{d} = -\frac{\kappa}{2}d - i \sum_{n=1,2} \bar{a}_n e^{-i\Delta_n t} \sum_{k=a,b} g_k \bar{a} (c_k + c_k^*) \\ \dot{c}_k = -\left(\frac{\gamma_k}{2} + i\omega_k\right) c_k - ig_k \sum_{n=1,2} (\bar{a}_n e^{-i\Delta_n t} d^*(t) + \bar{a}_n^* e^{+i\Delta_n t} d(t)) + \sqrt{\gamma_k}\eta_k \end{cases} \quad (5.17)$$

Note that here, and in subsequent expressions, $n = 1, 2$ indexes the two optical drives, while $k = a, b$ indexes the mechanical modes.

We now will move to the Fourier domain, as we typically do in the single-mode case, but taking special note of the fact that $\mathfrak{F}\{e^{-i\Delta t} f(t)\}[\omega] = \mathfrak{F}\{f(t)\}[\omega - \Delta]$. Doing so, our system becomes

$$\begin{cases} -i\omega d[\omega] = -\frac{\kappa}{2}d[\omega] - i \sum_{n=1,2} a_n \sum_{k=a,b} g_k (c_k[\omega - \Delta_n] + c_k^*[\omega - \Delta_n]) \\ -i\omega c_k[\omega] = -\left(\frac{\gamma_k}{2} + i\omega_k\right) c_k[\omega] - ig_k \sum_{n=1,2} (a_n d^*[\omega - \Delta_n] + a_n^* d[\omega + \Delta_n]) + \sqrt{\gamma_k}\eta_k[\omega] \end{cases} \quad (5.18)$$

Rearranging, and defining¹ a cavity susceptibility $\chi_c[\omega] = \left(\frac{\kappa}{2} - i\omega\right)^{-1}$, we have:

$$\begin{cases} d[\omega] = \chi_c[\omega] \left(-i \sum_{n=1,2} \sum_{k=a,b} g_k a_n (c_k[\omega - \Delta_n] + c_k^*[\omega - \Delta_n]) \right) \\ c_k[\omega] \left(\frac{\gamma_k}{2} - i(\omega - \omega_k) \right) = -ig_k \sum_{n=1,2} (a_n d^*[\omega - \Delta_n] + a_n^* d[\omega + \Delta_n]) + \sqrt{\gamma_k}\eta_k[\omega] \end{cases} \quad (5.19)$$

1. Note that this is not quite the usual definition, as Δ has been left out

To proceed, we will also need $d^*[\omega]$ and $c_k^*[\omega]$. We should be careful here, as conjugation of Fourier variables and susceptibilities must be done carefully, and there are some explicit time-dependent terms to be careful of. Namely, we'll use the fact that

$$\mathfrak{F} \{ e^{-i\Delta t} f(t) \} [\omega] = \mathfrak{F} \{ f(t) \} [\omega - \Delta] \quad (5.20)$$

but

$$\mathfrak{F} \left\{ (e^{-i\Delta t} f(t))^* \right\} [\omega] = \mathfrak{F} \{ e^{+i\Delta t} f^*(t) \} = \mathfrak{F} \{ f(t) \} [\omega + \Delta] \quad (5.21)$$

In the end, we have:

$$\left\{ \begin{array}{l} d[\omega] = -i\chi_c[\omega] \sum_{n=1,2} \sum_{k=a,b} g_k a_n (c_k[\omega - \Delta_n] + c_k^*[\omega - \Delta_n]) \\ d^*[\omega] = +i\chi_c[\omega] \sum_{n=1,2} \sum_{k=a,b} g_k a_n^* (c_k[\omega + \Delta_n] + c_k^*[\omega + \Delta_n]) \\ c_k[\omega] \left(\frac{\gamma_k}{2} - i(\omega - \omega_k) \right) = -ig_k \sum_{n=1,2} (a_n d^*[\omega - \Delta_n] + a_n^* d[\omega + \Delta_n]) + \sqrt{\gamma_k} \eta_k[\omega] \\ c_k^*[\omega] \left(\frac{\gamma_k}{2} - i(\omega + \omega_k) \right) = +ig_k \sum_{n=1,2} (a_n^* d[\omega + \Delta_n] + a_n d^*[\omega - \Delta_n]) + \sqrt{\gamma_k} \eta_k^*[\omega] \end{array} \right. \quad (5.22)$$

Substituting the optical expressions into the mechanical ones, we have:

$$\begin{aligned} c_k[\omega] \left(\frac{\gamma_k}{2} - i(\omega - \omega_k) \right) = & \\ & -ig_k \sum_{n=1,2} \left(ia_n \chi_c[\omega - \Delta_n] \sum_{m=1,2} \sum_{l=a,b} g_l a_m^* (c_l[\omega + \Delta_m - \Delta_n] + c_l^*[\omega + \Delta_m - \Delta_n]) \right) \end{aligned} \quad (5.23)$$

Simplifying and defining $\Delta_{nm} = \Delta_n - \Delta_m$, we have

$$c_k[\omega] \left(\frac{\gamma_k}{2} - i(\omega - \omega_k) \right) = \quad (5.24)$$

$$g_k \sum_{n=1,2} \sum_{m=1,2} \sum_{l=a,b} \left[(g_l a_n a_m^* \chi_c[\omega - \Delta_n] (c_l[\omega - \Delta_{nm}] + c_l^*[\omega - \Delta_{nm}]) \right. \quad (5.25)$$

$$\left. - g_l a_n^* a_m \chi_c[\omega + \Delta_n] (c_l[\omega + \Delta_{nm}] + c_l^*[\omega + \Delta_{nm}]) \right] \quad (5.26)$$

In the second term of the summation, exchange the indices n, m (since we are summing over both), and use the fact that $\Delta_{nm} = -\Delta_{mn}$ to arrive at

$$c_k[\omega] \left(\frac{\gamma_k}{2} - i(\omega - \omega_k) \right) = g_k \sum_{n=1,2} \sum_{m=1,2} \sum_{l=a,b} g_l a_n a_m^* (\chi_c[\omega - \Delta_n] - \chi_c[\omega + \Delta_n]) (c_l[\omega - \Delta_{nm}] + c_l^*[\omega - \Delta_{nm}]) \quad (5.27)$$

Now, defining $\tilde{\Sigma}_{nm} = a_n a_m^* (\chi_c[\omega - \Delta_n] - \chi_c[\omega + \Delta_m])$ (reminiscent of, though not equivalent to, the self-energy matrix element from Ch. 4), we can write:

$$c_k[\omega] \left(\frac{\gamma_k}{2} - i(\omega - \omega_k) \right) = g_k \sum_{n=1,2} \sum_{m=1,2} \sum_{l=a,b} g_l \tilde{\Sigma}_{nm}[\omega] (c_l[\omega - \Delta_{nm}] + c_l^*[\omega - \Delta_{nm}]) \quad (5.28)$$

While we have succeeded in reducing this expression to a fairly compact form, it can actually be helpful at this point to expand the two sums over optical indices, as we will be able to simplify some terms:

$$c_k[\omega] \left(\frac{\gamma_k}{2} - i(\omega - \omega_k) \right) = g_k \sum_{l=a,b} g_l \left[\tilde{\Sigma}_{11}[\omega] (c_l[\omega - \Delta_{11}] + c_l^*[\omega - \Delta_{11}]) \right. \quad (5.29)$$

$$\left. + \tilde{\Sigma}_{22}[\omega] (c_l[\omega - \Delta_{22}] + c_l^*[\omega - \Delta_{22}]) \right. \quad (5.30)$$

$$\left. + \tilde{\Sigma}_{12}[\omega] (c_l[\omega - \Delta_{12}] + c_l^*[\omega - \Delta_{12}]) \right. \quad (5.31)$$

$$\left. + \tilde{\Sigma}_{21}[\omega] (c_l[\omega - \Delta_{21}] + c_l^*[\omega - \Delta_{21}]) \right] \quad (5.32)$$

Now, we note the fact that $\Delta_{11} = \Delta_{22} = 0$ and $\Delta_{21} = -\Delta_{12}$, allowing us to rewrite:

$$c_k[\omega] \left(\frac{\gamma_k}{2} - i(\omega - \omega_k) \right) = g_k \sum_{l=a,b} g_l \left[(\tilde{\Sigma}_{11}[\omega] + \tilde{\Sigma}_{22}[\omega]) (c_l[\omega] + c_l^*[\omega]) \right. \quad (5.33)$$

$$\left. + \tilde{\Sigma}_{12}[\omega] (c_l[\omega - \Delta_{12}] + c_l^*[\omega - \Delta_{12}]) \right. \quad (5.34)$$

$$\left. + \tilde{\Sigma}_{21}[\omega] (c_l[\omega + \Delta_{12}] + c_l^*[\omega + \Delta_{12}]) \right] \quad (5.35)$$

Considering the first term in the square brackets, we can drop several of its components. First, as we do in the single-mode case, we can assume that the $c_l^*[\omega]$ terms will not drive $c_k[\omega]$, since they are rotating at $-\omega_a$ and $-\omega_b$, and $c_k[\omega]$ is rotating at either $+\omega_a$ or $+\omega_b$. Moreover, since ω_a is well-separated from ω_b , we can drop the cross-term and assume that $c_l[\omega]$ will only drive $c_k[\omega]$ if $k = l$. So, what we have left for this term is:

$$g_k^2 (\tilde{\Sigma}_{11}[\omega] + \tilde{\Sigma}_{22}[\omega]) c_k[\omega] \quad (5.36)$$

By inspection of our definition of $\tilde{\Sigma}$, we see that this term is simply the usual dynamical backaction one would expect from each optical drive independently – that is, the optical spring and damping from lasers detuned to Δ_1 and Δ_2 .

Now let's consider the second term:

$$g_k \sum_{l=a,b} g_l \tilde{\Sigma}_{12}[\omega] (c_l[\omega - \Delta_{12}] + c_l^*[\omega \Delta_{12}]) + \tilde{\Sigma}_{21}[\omega] (c_l[\omega + \Delta_{12}] + c_l^*[\omega + \Delta_{12}]) \quad (5.37)$$

For each mode ($k = a, b$), this expression presents 8 possible drive terms: 4 coming from the motion of the mode itself (with a frequency offset), and 4 coming from motion of the other mode, also shifted in frequency. We can eliminate many of these terms, by carefully considering where they have non-zero spectral content, and relying on the idea that these are high-Q modes, which only respond to drives near their resonant frequency.

Consider first the self-drive terms (the $l = k$ terms in the summation): Clearly, the $c_k[\omega \pm \Delta_{12}]$ terms will not resonantly drive $c_k[\omega]$. The other terms are of the form $c_k^*[\omega \pm \Delta_{12}]$. In principle, there could be resonant drive terms here, if $\Delta_{12} = 2\omega_k$, but this cannot be since $\omega_a \neq \omega_b$. Thus, all of these self-drive terms in Eq. 5.37 can be ignored – the sum over l need only be considered for $l \neq k$, with the following four terms remaining:

$$g_k g_l \left(\tilde{\Sigma}_{12}[\omega] (c_l[\omega - \Delta_{12}] + c_l^*[\omega - \Delta_{12}]) + \tilde{\Sigma}_{21}[\omega] (c_l[\omega + \Delta_{12}] + c_l^*[\omega + \Delta_{12}]) \right) \quad (5.38)$$

At this point, the relevancy of these terms will depend on the drive regime chosen. As was sketched in Fig. 5.1, we will choose to work in a regime in which our two optical drives are separated by approximately the difference in mechanical frequencies: $\Delta_{12} \approx \omega_b - \omega_a$. (Recall, $\omega_b > \omega_a$ and $\Delta_{12} > 0$). To see which terms are relevant, let's consider each mechanical mode separately, beginning with $k = a$:

$$c_a[\omega] \left(\frac{\gamma_a}{2} - i(\omega - \omega_a) \right) = g_a^2 (\tilde{\Sigma}_{11}[\omega] + \tilde{\Sigma}_{22}[\omega]) c_a[\omega] \quad (5.39)$$

$$+ g_a g_b \tilde{\Sigma}_{12}[\omega] (c_b[\omega - \Delta_{12}] + c_b^*[\omega - \Delta_{12}]) \quad (5.40)$$

$$+ g_a g_b \tilde{\Sigma}_{21}[\omega] (c_b[\omega + \Delta_{12}] + c_b^*[\omega + \Delta_{12}]) \quad (5.41)$$

Knowing that c_a is only sensitive to terms which are non-zero for $\omega \approx \omega_a$, we see that only the term $c_b[\omega + \Delta_{12}]$ will be relevant here, since $c_b[\omega_a + \Delta_{12}] \approx c_b[\omega_b] \neq 0$. The other three terms will be negligible, and we have just:

$$c_a[\omega] \left(\frac{\gamma_a}{2} - i(\omega - \omega_a) \right) = g_a^2 \left(\tilde{\Sigma}_{11}[\omega] + \tilde{\Sigma}_{22}[\omega] \right) c_a[\omega] + g_a g_b \tilde{\Sigma}_{21}[\omega] c_b[\omega + \Delta_{12}] \quad (5.42)$$

The same arguments for c_b lead us to:

$$c_b[\omega] \left(\frac{\gamma_b}{2} - i(\omega - \omega_b) \right) = g_b^2 \left(\tilde{\Sigma}_{11}[\omega] + \tilde{\Sigma}_{22}[\omega] \right) c_b[\omega] + g_a g_b \tilde{\Sigma}_{12}[\omega] c_a[\omega - \Delta_{12}] \quad (5.43)$$

As we usually do in the single-mode case, we are still able to eliminate the frequency dependence of Σ_{ij} , by again using the fact that c_k only oscillates near ω_k , and $\Sigma_{ij}[\omega]$ is constant on the scale of $\tilde{\gamma}_k$, so we need only evaluate it at that particular frequency (for $c_{a(b)}$, the $\tilde{\Sigma}$ terms should all be evaluated at $\omega_{a(b)}$).

From Eqs. 5.42 and 5.43, we can see the coupling that was mentioned previously: Whenever $\Delta_{12} \approx (\omega_b - \omega_a)$, there will be a coupling between modes c_a and c_b . (More precisely, the difference between Δ_{12} and $\omega_b - \omega_a$ should be smaller than the linewidth of the mode(s) in order to generate significant coupling.) Conceptually, one can understand this coupling by considering the radiation-pressure forces created by the optical drives and their sidebands. For example, the ω_1 laser will beat with the ω_b sideband of the ω_2 laser, at a frequency ω_a . At the same time, the ω_2 laser will beat the ω_a sideband of the ω_1 laser, at a frequency ω_b . Thus, the motion of $c_{a(b)}$ is mapped onto a force at $\omega_{b(a)}$, creating the desired coupling of the modes.

5.3 Optomechanical eigenvalues in presence of bichromatic drive

Instead of considering eigenvalues of this system directly, we will first rotate the frame of each mode by $\pm\Delta_{12}$:

$$c'_a(t) = e^{-\frac{i\Delta_{12}t}{2}} c_a(t) \quad (5.44)$$

$$c'_b(t) = e^{+\frac{i\Delta_{12}t}{2}} c_b(t) \quad (5.45)$$

In the Fourier domain, this transformation means:

$$c'_a[\omega] = c_a[\omega - \frac{\Delta_{12}}{2}] \quad (5.46)$$

$$c'_b[\omega] = c_b[\omega + \frac{\Delta_{12}}{2}] \quad (5.47)$$

So, our system of equations transforms to:

$$c_a[\omega] \left(\frac{\gamma_a}{2} - i(\omega - \omega_a) \right) = g_a^2 \left(\tilde{\Sigma}_{11}[\omega] + \tilde{\Sigma}_{22}[\omega] \right) c_a[\omega] + g_a g_b \tilde{\Sigma}_{21}[\omega] c_b[\omega + \Delta_{12}] \quad (5.48)$$

$$c_b[\omega] \left(\frac{\gamma_b}{2} - i(\omega - \omega_b) \right) = g_b^2 \left(\tilde{\Sigma}_{11}[\omega] + \tilde{\Sigma}_{22}[\omega] \right) c_b[\omega] + g_a g_b \tilde{\Sigma}_{12}[\omega] c_a[\omega - \Delta_{12}] \quad (5.49)$$

$$\begin{aligned} c_a[\omega - \frac{\Delta_{12}}{2}] \left(\frac{\gamma_a}{2} - i(\omega - \frac{\Delta_{12}}{2} - \omega_a) \right) &= g_a^2 \left(\tilde{\Sigma}_{11}[\omega_a] + \tilde{\Sigma}_{22}[\omega_a] \right) c_a[\omega - \frac{\Delta_{12}}{2}] \\ &\quad + g_a g_b \tilde{\Sigma}_{21}[\omega_a] c_b[\omega + \frac{\Delta_{12}}{2}] \end{aligned} \quad (5.50)$$

$$\begin{aligned} c_b[\omega + \frac{\Delta_{12}}{2}] \left(\frac{\gamma_b}{2} - i(\omega + \frac{\Delta_{12}}{2} - \omega_b) \right) &= g_b^2 \left(\tilde{\Sigma}_{11}[\omega_b] + \tilde{\Sigma}_{22}[\omega_b] \right) c_b[\omega + \frac{\Delta_{12}}{2}] \\ &\quad + g_a g_b \tilde{\Sigma}_{12}[\omega_b] c_a[\omega - \frac{\Delta_{12}}{2}] \end{aligned} \quad (5.51)$$

$$c'_a[\omega] \left(\frac{\gamma_a}{2} - i(\omega - \omega'_a) \right) = g_a^2 \left(\tilde{\Sigma}_{11}[\omega_a] + \tilde{\Sigma}_{22}[\omega_a] \right) c'_a[\omega] + g_a g_b \tilde{\Sigma}_{21}[\omega_a] c'_b[\omega] \quad (5.52)$$

$$c'_b[\omega] \left(\frac{\gamma_b}{2} - i(\omega - \omega'_b) \right) = g_b^2 \left(\tilde{\Sigma}_{11}[\omega_b] + \tilde{\Sigma}_{22}[\omega_b] \right) c'_b[\omega] + g_a g_b \tilde{\Sigma}_{12}[\omega_b] c'_a[\omega] \quad (5.53)$$

where we have now defined $\omega'_a = \omega_a + \frac{\Delta_{12}}{2}$ and $\omega'_b = \omega_b - \frac{\Delta_{12}}{2}$. We could also step back to the time-domain, as we did in the nearly-degenerate derivation, and write the evolution of these modes as:

$$i\dot{\bar{c}} = \mathbf{H}\bar{c} \quad (5.54)$$

where

$$\bar{c} = \begin{pmatrix} c'_a \\ c'_b \end{pmatrix} \quad (5.55)$$

$$\mathbf{H} = \begin{pmatrix} \omega'_a - i\frac{\gamma_a}{2} & 0 \\ 0 & \omega'_b - i\frac{\gamma_b}{2} \end{pmatrix} + \tilde{\Sigma} \quad (5.56)$$

$$\tilde{\Sigma} = \begin{pmatrix} g_a^2 \left(\tilde{\Sigma}_{11}[\omega_a] + \tilde{\Sigma}_{22}[\omega_a] \right) & g_a g_b \tilde{\Sigma}_{21}[\omega_a] \\ g_a g_b \tilde{\Sigma}_{12}[\omega_b] & g_b^2 \left(\tilde{\Sigma}_{11}[\omega_b] + \tilde{\Sigma}_{22}[\omega_b] \right) \end{pmatrix} \quad (5.57)$$

Here, we see that our new variables c'_a and c'_b describe a pair of coupled oscillators at frequencies ω'_a and ω'_b . Since $\Delta_{12} \approx (\omega_b - \omega_a)$, we know that $\omega'_a \approx \omega'_b$, and thus these modes are nearly-degenerate in this frame. More importantly, the effective frequencies of these modes $\omega'_{a,b}$ are dictated by Δ_{12} , which is entirely under our control. Thus, we have arrived at a pair of coupled oscillator equations in which we can not only tune the coupling (by tuning the various elements of $\tilde{\Sigma}$), but also the bare mode splitting (by changing Δ_{12}). In this way, we actually have more control over the system than in the previous experiment. Let's now consider how the eigenvalues of this system depend on our control parameters, and whether it is possible to access an EP.

The first point we must consider is what control parameters we wish to use. Each optical drive has a power (P_1, P_2) and a detuning (Δ_1, Δ_2), giving us a total of 4 ‘knobs’ with which to tune the system. Since the parameter space which breaks an EP degeneracy is 2-dimensional, it is sufficient to consider 2 linear combinations of these control parameters, and to look for an EP in this control space. There are many such combinations, but we will consider one which has the benefit of being particularly easy to implement experimentally. Specifically, we will choose to jointly manipulate the power/detuning of both optical drives, as follows:

Let $P_1 = P_2 = \bar{P}$ be one control parameter (i.e. change both beam powers jointly). Let $\Delta_1 = \bar{\Delta} + \Delta_{12}/2$ and $\Delta_2 = \bar{\Delta} - \Delta_{12}/2$, with the control parameter being $\bar{\Delta}$ (the mean detuning of the two beams). Δ_{12} will be chosen to be nearly (though not exactly) equal to $\omega_b - \omega_a$. Specifically, let $\Delta_{12} = (\omega_b - \omega_a) + \delta$. Thus, one finds that in the reference frame of Eq. 5.56, the ‘bare’ modes are split by δ . If we choose this δ to be on the order of 100 Hz, then this scheme is now qualitatively similar to our previous experiment: We have two ‘bare’ modes who are separated by a small amount δ , and we can use the optical spring/damping of the control beams to ‘bridge’ this gap and create a degeneracy.

Having established these control parameters, we can examine the eigenvalues of \mathbf{H} as defined in Eq. 5.56, plotting the real and imaginary parts as we did before. As a concrete example, we’ll consider the $\{2, 1\}$ and $\{2, 2\}$ modes of the membrane, with frequencies $\omega_a/(2\pi) = 557.473$ kHz, $\omega_b/(2\pi) = 705.164$ kHz, linewidths $\gamma_a/(2\pi) = 0.39$ Hz, $\gamma_b/(2\pi) = 0.38$ Hz, and coupling rates $g_a/(2\pi) = 2.08$ Hz, $g_b/(2\pi) = 2.09$ Hz. (These parameters come from fits detailed later). Fig. 5.2 shows the calculated real/imaginary parts of the eigenvalues of \mathbf{H} as defined in Eq. 5.56. Note that the sheets possess the same topology seen in the previous chapter, with both the frequencies and linewidths becoming degenerate for $\bar{\Delta} = -15$ kHz, $\bar{P} = 4.7$ μ W.

At this point, it is worth reemphasizing that the mechanical frequencies here are considered in the shifted frame defined earlier (shifting the lower/upper-frequency mode by $\pm \Delta_{12}$), which is approximately near $\frac{\omega_a + \omega_b}{2}$. Since the modes of the mechanical modes themselves remain near ω_a and ω_b , and do not become degenerate in the traditional sense, we refer to the degeneracy of Fig. 5.2 as a virtual exceptional point. Note that this is not




Figure 5.2: Calculated eigenvalue sheets for the effective Hamiltonian defined in Eq. 5.56. Note that the frequency/linewidth sheets are plotted from different perspectives, for clarity. The mechanical frequencies plotted are relative to the mean mechanical frequency $((\omega_a + \omega_b)/2)$

the only coordinate frame we could have considered. For example, rather than shifting each mode up/down by half the frequency difference, we also could have shifted one of the modes by the full splitting ($\omega'_a = \omega_a + \Delta_{12}$, $\omega'_b = \omega_b$), which would simply apply an overall shift to the sheets plotted in Fig. 5.2. Despite this not being a traditional degeneracy, we'll see in upcoming sections that the topology of the eigenvalues in this reference frame still leads to the same behavior seen previously, in terms of adiabatic energy transfer.

5.4 Experimental Implementation

The experimental setup from Ch. 2 is easily modified to allow for multiple control beams. Previously, the control beam power/detuning was controlled by sending it through an AOM (AOM2), which shifted its frequency by 80 MHz. By simply adding an additional RF drive tone to this AOM at $80 \text{ MHz} \pm \Delta_{12}$, the AOM will now output two beams separated in frequency by Δ_{12} . These will be the control beams, and their powers and frequencies are tunable by modifying the RF drive tones, just as before.

The circuit used to accomplish this is illustrated in Fig. 5.3. The HF2 is used as the source of both RF drive tones, via its ability to output multiple RF signals on a single voltage channel. The bandwidth of the HF2 is 50 MHz, so these tones are output near 20 MHz, then mixed up to the 80 MHz frequency necessary for driving AOM2. The HF2 simultaneously uses these same internal oscillators for demodulation of the heterodyne signal, which has

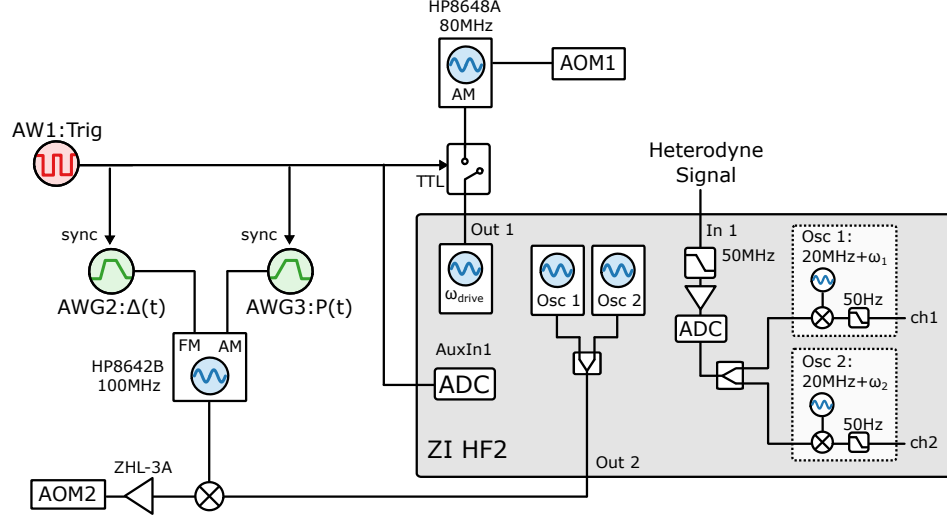


Figure 5.3: Electrical schematic for controlling/measuring non-degenerate modes. Note that oscillators Osc 1 and Osc 2 (at frequencies $20\text{ MHz} + \omega_1$ and $20\text{ MHz} + \omega_2$) are used both for demodulation and for driving the control AOM (after being summed in the HF2 and mixed-up externally). Modulation voltages applied to the 100 MHz local oscillator translate to power/frequency modulation of the control beams

similarly been mixed down to near 20 MHz.

In the previous experiment, we applied modulation voltages to the AM/FM ports of the 80 MHz signal generator driving the control AOM. Now, we apply this modulation to the 100 MHz local oscillator which up-mixes the tones from the HF2. Thus, the net result is a pair of tones near 80 MHz whose frequency splitting can be set digitally by the HF2, and whose power/frequency can be modified by sending modulation voltages into the 100 MHz signal generator.

Otherwise, the setup is similar to before: we use separate demodulation channels near ω_a and ω_b to record the energy of each mechanical mode. We still optically drive the membrane for initial excitation (or response measurements), but now we apply the necessary modulation to the measurement beam, instead of the control beam².

5.5 Measuring Mechanical Eigenvalues

As in the nearly-degenerate EP experiment, we first wish to confirm the expected eigenvalue topology by measuring the mode frequencies and linewidths across a grid of control

². With multiple control beams now, it is simply more straightforward to modulate only the measurement beam for this task.




Figure 5.4: Driven response measurements in the original and shifted frames. The upper plot shows the original driven sweeps, with peaks near ω_a and ω_b . The lower plot shows the same data after it has been shifted into the common frame, with peaks near ω'_a and ω'_b .

parameters. The control beams are created as described in the previous section, set to particular values of \bar{P} and $\bar{\Delta}$, and then a driven response measurement is performed. The only significant difference now is that the drive tone (and the mechanical response) must be applied at two separate frequencies: ω_a and ω_b . The translation into the nearly-degenerate frame will occur in post-processing. This process is illustrated in Fig. 5.4. First, we perform two separate driven response measurements, in which the frequency of an optical force is swept near ω_a and the measurement signal is demodulated ω_a away from the heterodyne carrier. Then an optical force is swept near ω_b , and the measurement signal is demodulated accordingly. These two results are shown in Fig. 5.4 (upper panel). Then, the frequencies of these measurements are shifted in post-processing, by adding (subtracting) Δ_{12} from the frequency axis of the mode $a(b)$ data. These two shifted data sets are plotted on the same axes in Fig. 5.4 (lower panel). It is in this new frequency space that we fit the data to extract the frequencies and linewidths.

5.5.1 Fitting Overlapping Modes

In the data discussed above (and plotted in Fig. 5.4), each mode appears as a well-defined peak, which we might fit to a generic Lorentzian as we did in Ch. 4. However, such is not always the case, due to the coupling mechanism described previously. Recalling the




Figure 5.5: Driven response measurement showing coupled motion. The blue data is from a driven response measurement from near ω_b (705kHz), while the red data is from a driven response measurement near ω_a (557kHz). Once they are moved into the shifted reference frame, we see that the extra ‘bump’ is from mode b motion driving mode a .

previous discussion, we know that the motion of each mode beats with one of the optical drive tones, resulting in a force on the other mode. Thus, for example, the motion of mode b exerts a force on mode a . Since the motion of mode b is a Lorentzian peak, there is a Lorentzian force profile acting on mode a . This Lorentzian force will be near to ω_a (by design, since $\omega'_a \approx \omega'_b$), but it will not in general be aligned with ω_a . The result is that the motion of mode a will reflect two Lorentzian contributions: its own response to a white thermal noise drive (producing a peak at ω_a), and its response to the Lorentzian drive of mode b .

To fit data like this, it is important to recognize that one should not fit this response to a pair of Lorentzians with arbitrary frequencies and linewidths. For example, when fitting the data that was originally near ω_a : the extra ‘bump’ caused by the motion of the ω_b mode should be constrained by the frequency and linewidth of that b mode, which is known from the other data set (collected near ω_b). This is illustrated in Fig. 5.5. We see that the data collected near ω_b seems to reflect two Lorentzian resonances, the smaller one arising from the motion of mode a . This is clear once we move both data sets to the common frequency frame, and note that the bump in the data from ω_b aligns perfectly with the peak from the motion near ω_a .

The appropriate procedure for fitting this data is then to move both data sets into the common frame, then simultaneously fit both data sets with a pair of Lorentzians which are constrained to have the same frequencies and linewidths. While the frequencies/linewidths

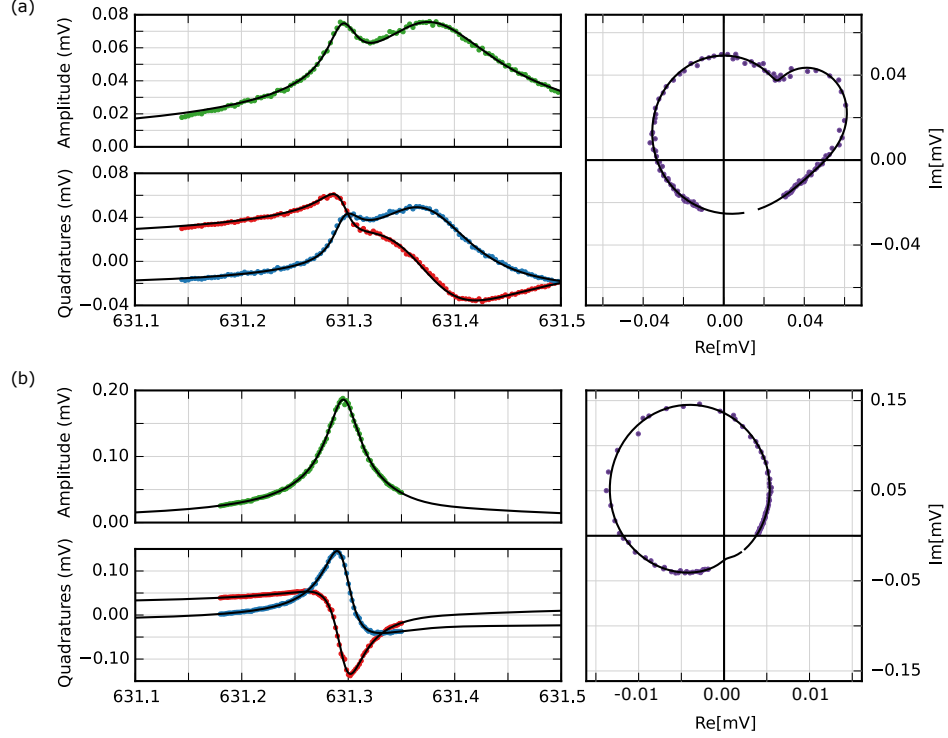


Figure 5.6: Sample driven response measurements and fits, as described in the main text. The top data (a) came from a response measurement near ω_a , and the bottom data (b) came from a response measurement near ω_b . Both data have been shifted into the intermediate reference frame. Each response measurement provides a set of complex values as a function of frequency, which have been plotted as quadratures, amplitudes, and parametrically in the complex plane. Both data sets are fit simultaneously, and constrained to use the same frequencies and linewidths ($\omega'_\alpha, \omega'_\beta, \gamma_\alpha, \gamma_\beta$)

of the modes are constrained to be shared between data sets, each Lorentzian is allowed an independent complex amplitude. A more-complete model would further constrain some of these fit parameters. For example, the size of the transduced-motion bump near ω_b should be directly linked to the size of the peak at ω_a . A fully-accurate model could be found by calculating the response of the two-mode system to an applied optical drive. For our purposes, we simply wish to extract the eigenvalues, so we take a more generic approach, constraining common frequencies/linewidths, but allowing unconstrained amplitudes. The model to which the data is fit is then:

$$\text{Data from near } \omega_a: \quad b_1 + \frac{u_1}{\frac{\gamma_\alpha}{2} - i(\omega - \omega_\alpha)} + \frac{v_1}{\frac{\gamma_\beta}{2} - i(\omega - \omega_\beta)} \quad (5.58)$$

$$\text{Data from near } \omega_b: \quad b_2 + \frac{u_2}{\frac{\gamma_\alpha}{2} - i(\omega - \omega_\alpha)} + \frac{v_2}{\frac{\gamma_\beta}{2} - i(\omega - \omega_\beta)} \quad (5.59)$$

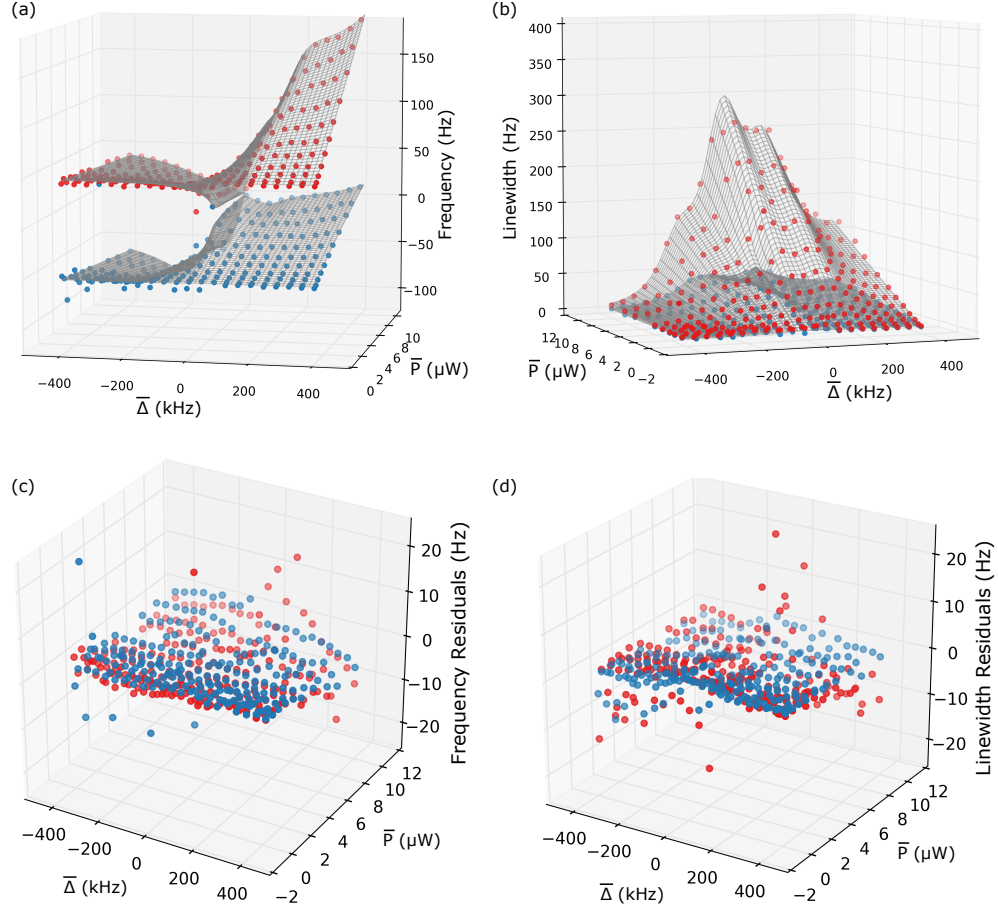


Figure 5.7: (a),(b) Measured eigenvalues, plotted as frequency/linewidth sheets. The gray mesh sheets are the result of fits to the theoretical model, as described in the main text. (c),(d) Residuals from the fit shown in (a) and (b), showing good agreement with the model (most residuals are within ± 3 Hz).

Note that we use subscripts α, β to refer to the perturbed eigenvalues in the common frequency frame (i.e. the eigenvalues of \mathbf{H} as defined in Eq. 5.56). Note that, as before, we are fitting the full complex response of the modes, and thus both the backgrounds (b_1, b_2) and amplitudes (u_1, v_1, u_2, v_2) are allowed to be complex variables. (In practice, we split the complex response data into real/imaginary parts, and fit to the real/imaginary parts of the expressions above. This is the same process as was used in the previous chapter). An example of driven response data and fits to the above expressions are shown in Fig. 5.6.

Repeating this process for various values of \bar{P} and $\bar{\Delta}$ allows us to accumulate a set of perturbed frequencies and linewidths $\omega_{\alpha, \beta}(\bar{P}, \bar{\Delta})$ and $\gamma_{\alpha, \beta}(\bar{P}, \bar{\Delta})$. In the previous experiment, recall that after collecting a set of mechanical eigenvalues, we fit a subset of that




Figure 5.8: Measured eigenvalue differences, calculated from the data shown in Fig. 5.7. The gray mesh is a theoretical calculation, based on the same fit as in Fig. 5.7. The fact that the gray mesh does not go exactly to zero is purely a plotting artifact.

data to the full optomechanical model, in order to estimate the values of certain system parameters (the intrinsic mechanical frequencies/linewidths and the optomechanical coupling rates g_a and g_b). Here, we can follow the same procedure to estimate the same system parameters. Previously, the eigenvalue data/fits were plotted in the complex plane (Fig. 4.6). In the non-degenerate data, the eigenvalue contours in the complex plane are not so illustrative, so we will only work with the three-dimensional eigenvalue sheet visualization. Fig. 5.7 shows the eigenvalue data collected from driven response measurements, along with sheets calculated from the model of Sec. 5.3. These theoretical sheets reflect the result of a least-squares fit, in which the system parameters $(\omega_a, \omega_b, g_a, g_b)$ are allowed to vary³. To make clear the agreement between the data points and the theoretical sheets, Fig. 5.7 also shows the residuals from the fits, which indicate that the data seems to be well-described by the model from Sec. 5.3.

The topology of the eigenvalue sheets in this experiment is such that it can be somewhat

3. It is possible, though not necessary to include γ_a and γ_b as fit parameters. However, they are more accurately estimated by ringdown measurements, and do not significantly affect this least-squares fit. I do also allow for a systematic offset to $\bar{\Delta}$ – it is common for us to include a detuning offset as a fit parameter, due to our inability to accurately estimate our absolute detuning during a measurement. Relative detuning is well-defined, so a single overall offset per measurement is sufficient.




Figure 5.9: Example of adiabatic energy transfer. (a) The control loop in $\{\bar{P}, \bar{\Delta}\}$ space enclosing the VEP. The black square marks the beginning/end of the loop. (b) Measured energy of the modes at ω_a (red) and ω_b (blue) as a function of time. As in the previous experiment, the data during the control loop does not accurately reflect the mode energy.

difficult to visualize the EP in this way, so in Fig. 5.8 we also include a plot of the measured and calculated eigenvalue difference: $|(\omega_\alpha + i\gamma_\alpha) - (\omega_\beta + i\gamma_\beta)|$, which should (and does) vanish exactly at the EP.

5.6 Adiabatic Energy Transfer near a VEP

Having demonstrated the existence of a VEP in our system, we can now attempt to achieve the same sort of energy transfer via closed adiabatic cycles that we achieved in the previous chapter. We’ve already discussed the major experimental differences (slight as they were), so we can now simply attempt to replicate the same experiments. First, we’ll consider a ‘large’, ‘slow’ control loop, and confirm that energy transfer is at all possible. Then, we’ll systematically vary the size, speed, initial condition, and loop direction, to see if the VEP behaves similarly to the EP. As before, we will use numerical integration of the equation of motion (Eq. 5.54) to make predictions about how the energy transfer should behave.

First, the proof-of-principle demonstration: We consider a control loop in which $\bar{\Delta}$ is varied between -604 kHz and 376 kHz and \bar{P} is varied between 0.0 μ W and 8.3 μ W, as illustrated in Fig. 5.9(a). The time for the loop (τ) is chosen to be 27.2 ms. We apply a drive to the a mode before the loop, initializing the system approximately in the state




Figure 5.10: Measurement results in which the control loop geometry is varied. The control loop is systematically extended along each axis, demonstrating that energy transfer only occurs once the loop is large enough to enclose the VEP.

$(1 \ 0)^T$. The lock-in amplifier records the amplitude of each mode before and after the loop, and we average over many measurements as before, to produce the data plotted in Fig. 5.9(b). As before, the data does not accurately reflect the mode amplitude during the loop, but is accurate before and after, and shows the energy being swapped from the a mode to the b mode.

We now maintain the same initialization and loop direction, but systematically vary the loop geometry to confirm that the energy transfer is dependent on the topology of the sheets (specifically, on inclusion of the VEP within the control loop). Fig. 5.10 show the energy transfer efficiency varying with loop geometry as the control loop is stretched along the $\bar{\Delta}$ and \bar{P} axes, respectively. The black lines are the result of numerical simulation. As in the previous experiment, we see the energy transfer transition from near-zero to near-unity as the control loop moves from non-inclusion of the VEP to inclusion of the VEP. The intermediate behavior is complicated by non-adiabatic behavior when the loop passes very near the EP, but as before, we still find agreement with the numerical prediction.




Figure 5.11: Measurement results in which the loop time is varied, for both initial conditions and loop directions. For a given direction, it is possible to transfer energy from one particular initial mode, due to the relative gain/loss in the system. These data also confirm the asymmetry of this behavior with respect to loop direction. The black lines are the result of numerically integrating the equation of motion defined in Eq. 5.54

Now, we look for the same behavior with respect to adiabaticity, as well as checking for the asymmetry with respect to initialization and loop direction. Fig. 5.11 shows the transfer efficiency varying with loop time for both initializations and both loop directions. As before, we see no energy transfer in the sudden-perturbation limit, and near-unity energy transfer in the adiabatic limit. The intermediate-time behavior is more complicated than before, showing large variations in energy transfer. As before, we attribute this to an interplay between the adiabatic transport along one eigenmode and dynamical behavior arising from the presence of a time-dependent coupling. These times are essentially outside the ‘adiabatic’ regime we can understand intuitively, and the best we can do here is to consult the numerical simulations, with which we find good agreement.

Chapter 6

Conclusion and Outlook

Exceptional points occupy an appealing corner of physics, particularly for an experimentalist: a remarkably simple mathematical basis, describing a variety of physical systems, which nonetheless yields a rich set of emergent phenomena. Moreover, it is a field whose theoretical underpinnings were only recently laid out, and which has only seen sparse experimental progress over the last few decades. Through this thesis, I have hopefully presented a convincing case for optomechanics as a powerful platform for exploring these complex degeneracies, and in particular their associated dynamical phenomena. This strength arises from some rather fortunate timing: as theorists teased out the interesting behavior arising from time-dependent perturbations near complex degeneracies, the field of optomechanics was reaching maturity as a system for studying coupled oscillator dynamics. It is precisely this strength which we were able to exploit in this work, via a cavity optomechanical system which offers high-quality mechanical oscillators with a strong, tunable coupling. The work in Ch. 4 presents a first demonstration of this capability, with two key results: first, that we can create an EP in our system (and confirm its eigenvalue topology with great precision), and second, that we can study the dynamical behavior of the system, thanks to the easily-tunable optical coupling. Subsequently, in Ch. 5, a modified optical scheme was introduced, greatly expanding our choice of modes and coupling parameters for studying EPs. This scheme will allow greater flexibility in exploring more complex EP phenomena, as well as enabling EP studies in different classes of optomechanical systems, which may lack the near-degeneracy required in Ch. 4.

Having established itself as a platform for studying complex degeneracies, our optomechanical system is well-poised to probe various features and phenomena associated with EPs, including:

- **EP Sensing:** At an EP, the real/imaginary eigenvalues (frequencies/linewidths in our system) possess a square-root dependence on the coupling parameter, due to the branch point nature of the EP. This implies that by monitoring, for instance, the frequencies of the system, one is highly sensitive to perturbations in the coupling parameter. Exploiting this behavior for sensing applications has been theoretically proposed [44–46], and recently experimentally demonstrated in optical systems [47, 48]. (Note that one must be careful to consider the deleterious effects of adding dissipation to a system to create an EP – this may in principle work against the sensitivity offered by the EP). Given the flexibility of our system, it should be possible to explore such sensitivity, perhaps as a transducer of the quantum fluctuations of the cavity light.
- **Higher-order Degeneracies:** Higher-order EPs (the eigenvalue coalescence of more than two modes) has been theoretically studied [49] and very recently demonstrated experimentally [48, 50]. These degeneracies offer the potential to explore even greater sensing capabilities (in terms of a parameter perturbation ϵ , an n -th order EP offers an $\epsilon^{1/n}$ scaling of the eigenvalues), as well as explore more complex eigenvalue topology. There are even recent theoretical proposals [51] offering specific optomechanical advantages by exploiting higher-order exceptional points. Our system should be capable of exploring such higher-order EPs, in principle via multiple routes. First, just as we used a pair of nearly-degenerate mechanical modes as the basis of an EP in Ch. 4, our membrane naturally offers nearly-degenerate triplets (for instance, the $(i, j) = (1, 7), (7, 1)$, and $(5, 5)$ modes, in the language of Ch. 1). However, one would likely have even greater flexibility by employing the techniques of Ch. 5. One could, for instance, use one optical drive to create an EP as in Ch. 4, then add a second optical drive which brings a third mechanical mode into ‘effective’ degeneracy with the other two. Alternatively, one could employ three drives, bringing into degeneracy three arbitrary modes.

- **Additional dynamical behavior:** In the work presented here, we probed several of the core features associated with dynamical behavior near an EP, specifically the adiabatic and non-adiabatic evolution that occurs along the gain/loss eigenvalue surfaces. However, more subtle effects can occur, such as the very-long-time dominance of the gain mode mentioned in Appendix A. Other non-intuitive effects include the existence of certain adiabatic trajectories which predominantly follow the lossy eigenmode, via continuous destructive interference with a small ‘seed’ amplitude in the gain mode. These effects should be open to exploration using our system.

The preceding list is meant only to highlight a few possible avenues that might be explored using our system. The physics of exceptional points is still very much an active field, with rapid growth in both theoretical and experimental studies. This will certainly open even more topics, for which our optomechanical system will continue to be an ideal testbed. I hope that this work has provided a clear introduction to the field, and serves as a solid foundation on which to build future explorations.

Appendix A

Dominance of Gain Mode Beyond Adiabatic Regime

In Ch. 3, it was stated that a brief propagation along the lossy eigenmode can be sustained, since the growth in the gain mode will not have time to sufficiently dominate the component in the loss mode. (Recall from Fig. 3.6 that the adiabatic trajectory finished on a short section of blue sheet, after propagating mostly on the orange.) One might wonder whether, by perturbing the system very slowly, one would allow time for the gain mode to grow and cause a jump from the short blue section. Naïvely, one might suspect that this jump would never happen, since going more slowly means that there is additional time spent in the gain mode, which builds up an even larger component in that instantaneous eigenmode. But, [20] shows that, even when propagating along the gain mode, the system reaches a fixed ratio of c_+/c_- . That is, while propagating along the gain mode, the gain mode component is not growing exponentially, relative to the loss mode component. So, returning to our example, if we slow the loop down significantly, the relative projection in the gain mode stays roughly the same, but we allow much more time for a non-adiabatic jump to occur during the propagation along the blue mode. Recall that, for the perturbation in this example, a loop time of $\tau = 10$ was sufficient for adiabatic evolution. If we instead consider something like $\tau = 3.2 \times 10^4$, we find the trajectory shown in Fig. A.1, showing a non-adiabatic jump during this final section.




Figure A.1: System trajectory under a very slow perturbation ($\tau = 3.2 \times 10^4$, i.e. two orders of magnitude slower than the usual adiabatic limit). (a) Projection of $\bar{\lambda}$ on the real and imaginary sheets. (b) Real projection of $\bar{\lambda}$, along with the real eigenvalues throughout the control loop. The color of the eigenvalue indicates whether it has the larger (orange) or smaller (blue) imaginary part (the same convention used in the sheets).

This example implies that the relative transfer efficiency, which had been near unity for this path, now drops to zero. However, if we consider the opposite initial condition (which previously did not transfer energy), we find that it now experiences two non-adiabatic jumps. Thus, there is still energy transfer possible, but energy now preferentially flows to a different mode than before.

Fig. A.2 shows the transfer efficiency for both initial conditions and loop directions in this very-long-time limit. We see that qualitatively, most of the behavior we saw before still holds: there is still asymmetric energy transfer in the very-long-time limit, but both perturbation directions transfer energy to the same mode. In other words, the gain mode is always preferred in the very-long-time limit.

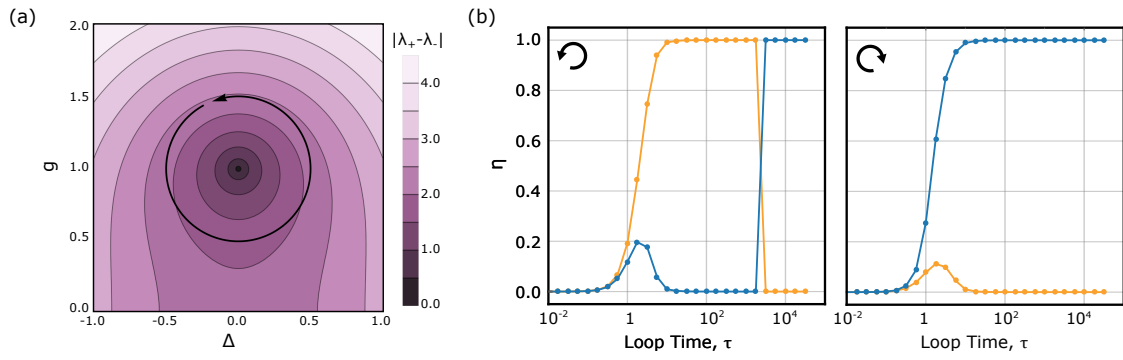


Figure A.2: Energy transfer dependence on initial condition, perturbation speed, and perturbation direction. (a) Illustration of control loop (counter-clockwise shown), overlaid on a contour plot of the eigenvalue difference, to illustrate the location of the EP. (b) Energy transfer as a function of loop duration, given counter-clockwise (left) and clockwise (right) perturbations, applied to a system initialized in the lower-frequency (orange) or higher-frequency (blue) mode. Note the switch in counter-clockwise transfer direction for $\tau > 4 \times 10^3$

Bibliography

- [1] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, *Reviews of Modern Physics* **86**, 1391 (2014).
- [2] C. Fabry and A. Perot, *Ann. Chim. Phys.* **16** (1899).
- [3] M. Underwood, D. Mason, D. Lee, H. Xu, L. Jiang, A. B. Shkarin, K. Børkje, S. M. Girvin, and J. G. E. Harris, *Physical Review A - Atomic, Molecular, and Optical Physics* **92**, 1 (2015).
- [4] J. C. Sankey, a. M. Jayich, B. M. Zwickl, C. Yang, and J. G. E. Harris, *ICAP 2008* (2009).
- [5] M. Underwood, Cryogenic Optomechanics with a Silicon Nitride Membrane, PhD thesis, Yale University, 2016.
- [6] B. M. Zwickl, W. E. Shanks, A. M. Jayich, C. Yang, A. C. Bleszynski Jayich, J. D. Thompson, and J. G. E. Harris, *Applied Physics Letters* **92**, 103125 (2008).
- [7] J. D. Thompson, B. M. Zwickl, A. M. Jayich, F. Marquardt, S. M. Girvin, and J. G. E. Harris, *Nature* **452**, 72 (2008).
- [8] A. M. Jayich, J. C. Sankey, B. M. Zwickl, C. Yang, J. D. Thompson, S. M. Girvin, A. A. Clerk, F. Marquardt, and J. G. E. Harris, *New Journal of Physics* **10** (2008).
- [9] J. C. Sankey, C. Yang, B. M. Zwickl, A. M. Jayich, and J. G. E. Harris, *Nature Physics* **6**, 707 (2010).

- [10] A. M. Jayich, J. C. Sankey, K. Borkje, D. Lee, C. Yang, M. Underwood, L. Childress, A. Petrenko, S. M. Girvin, and J. G. E. Harris, *New Journal of Physics* **14** (2012).
- [11] P. L. Yu, K. Cicak, N. S. Kampel, Y. Tsaturyan, T. P. Purdy, R. W. Simmonds, and C. A. Regal, *Applied Physics Letters* **104** (2014).
- [12] Y. Tsaturyan, A. Barg, E. S. Polzik, and A. Schliesser, *Nature Nanotechnology* (2017), Advance Online Publication.
- [13] R. W. Andrews, R. W. Peterson, T. P. Purdy, K. Cicak, R. W. Simmonds, C. A. Regal, and K. W. Lehnert, *Nature Physics* **10**, 321 (2014).
- [14] D. Lee, M. Underwood, D. Mason, A. B. Shkarin, S. W. Hoch, and J. G. E. Harris, *Nature Communications* **6**, 6232 (2015).
- [15] M. Born and V. A. Fock, *Zeitschrift für Physik* **51**, 165 (1928).
- [16] M. V. Berry, *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **392**, 45 (1984).
- [17] S. Pancharatnam, *Proceedings of the Indian Academy of Sciences - Section A* **42**, 235 (1955).
- [18] R. Uzdin, A. Mailybaev, and N. Moiseyev, *Journal of Physics A: Mathematical and Theoretical* **44**, 435302 (2011).
- [19] M. V. Berry and R. Uzdin, *Journal of Physics A: Mathematical and Theoretical* **44**, 435303 (2011).
- [20] T. J. Milburn, J. Doppler, C. A. Holmes, S. Portolan, S. Rotter, and P. Rabl, *Physical Review A* **92**, 052124 (2015).
- [21] J. C. Garrison and E. M. Wright, *Physics Letters A* **128**, 177 (1988).
- [22] C. Miniatura, C. Sire, J. Baudon, and J. Beuissard, *Europhys. Lett* **13**, 199 (1990).
- [23] G. Nenciu and G. Rasche, *J. Phys. A: Math. Gen. Phys. A Math. Gen* **25**, 5741 (1992).

- [24] M. Pont, R. M. Potvliege, R. Shakeshaft, and P. H. G. Smith, *Physical Review A* **46**, 555 (1992).
- [25] O. Latinne, N. J. Kylstra, M. Dorr, J. Purvis, M. Terao-Dunseath, C. J. Joachain, P. G. Burke, and C. J. Noble, *Physical Review Letters* **74**, 46 (1995).
- [26] R. Lefebvre, O. Atabek, M. Sindelka, and N. Moiseyev, *Physical Review Letters* **103** (2009).
- [27] O. Atabek, R. Lefebvre, M. Lepers, A. Jaouadi, O. Dulieu, and V. Kokoouline, *Physical Review Letters* **106** (2011).
- [28] H. Cartarius, J. Main, and G. Wunner, *Physical Review Letters* **99** (2007).
- [29] T. Kato, Perturbation Theory for Linear Operators Perturbation Theory for Linear Operators (Springer-Verlag, Berlin, New York, 1995).
- [30] W. D. Heiss, *J. Phys. A: Math. Theor.* *J. Phys. A: Math. Theor* **45**, 444016 (2012).
- [31] W. D. Heiss and A. L. Sannino, *Journal of Physics A* **23**, 1167 (1990).
- [32] W. D. Heiss and S. Radu, *Physical Review E* **52**, 4762 (1995).
- [33] W. D. Heiss, *Physical Review E* **61**, 929 (2000).
- [34] W. D. Heiss and H. L. Harney, *Eur. Phys. J. D* **17**, 149 (2001).
- [35] C. Dembowski, H.-D. Gräf, H. L. Harney, A. Heine, W. D. Heiss, H. Rehfeld, and A. Richter, *Physical Review Letters* **86**, 787 (2001).
- [36] B. Zhen, C. W. Hsu, Y. Igarashi, L. Lu, I. Kaminer, A. Pick, S.-L. Chua, J. D. Joannopoulos, and M. Soljačić, *Nature* **525**, 354 (2015).
- [37] S.-B. Lee, J. Yang, S. Moon, S.-Y. Lee, J.-B. Shim, S. W. Kim, J.-H. Lee, and K. An, *Physical Review Letters* **103**, 134101 (2009).
- [38] T. Gao, E. Estrecho, K. Y. Bliokh, T. C. H. Liew, M. D. Fraser, S. Brodbeck, M. Kamp, C. Schneider, S. Höfling, Y. Yamamoto, F. Nori, Y. S. Kivshar, A. G. Truscott, R. G. Dall, and E. A. Ostrovskaya, *Nature* **526**, 554 (2015).

- [39] T. Stehmann, W. D. Heiss, and F. G. Scholtz, *Journal of Physics A: Mathematical and General* **37**, 7813 (2004).
- [40] C. M. Bender and S. Boettcher, *Physical Review Letters* **80**, 5243 (1998).
- [41] R. El-Ganainy, K. G. Makris, D. N. Christodoulides, and Z. H. Musslimani, *Optics letters* **32**, 2632 (2007).
- [42] B. Peng, S. K. Ozdemir, S. Rotter, H. Yilmaz, and M. Liertzer, *Science* **346**, 328 (2014).
- [43] M. Brandstetter, M. Liertzer, C. Deutsch, P. Klang, J. Schöberl, H. E. Türeci, G. Strasser, K. Unterrainer, and S. Rotter, *Nature Communications* **5**, 4034 (2014).
- [44] Z.-P. Liu, J. Zhang, a. K. Özdemir, B. Peng, H. Jing, X.-Y. Lü, C.-W. Li, L. Yang, F. Nori, and Y.-x. Liu, *Physical Review Letters* **117**, 110802 (2016).
- [45] J. Wiersig, *Physical Review Letters* **112**, 203901 (2014).
- [46] J. Wiersig, *Physical Review A* **93**, 033809 (2016).
- [47] W. Chen, a. Kaya Özdemir, G. Zhao, J. Wiersig, and L. Yang, *Nature* **548**, 192 (2017).
- [48] H. Hodaei, A. U. Hassan, S. Wittek, H. Garcia-Gracia, R. El-Ganainy, D. N. Christodoulides, and M. Khajavikhan, *Nature* **548**, 187 (2017).
- [49] G. Demange and E.-M. Graefe, *J. Phys. A: Math. Theor.* **45**, 25303 (2012).
- [50] K. Ding, G. Ma, M. Xiao, Z. Q. Zhang, and C. T. Chan, *Physical Review X* **6** (2016), 1509.06886.
- [51] H. Jing, S. K. Özdemir, H. Lö, and F. Nori, *Scientific Reports* **7** (2017), 1609.01845.