Towards Analog Computing: An Electronically Reconfigurable Nanoelectromechanical Oscillator Network

Thesis by
Scott Thomas Philpott Habermehl

In Partial Fulfillment of the Requirements for the
Degree of
Doctorate of Philosophy in Applied Physics

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2026
Defended February 5th, 2026

© 2026

Scott Thomas Philpott Habermehl
ORCID: 0009-0006-2970-4580

All rights reserved

In memory of Dr. Scott Dwight Habermehl

ACKNOWLEDGEMENTS

I would first and foremost like to thank my advisor, Prof. Michael Roukes, for his support, patience, and inspiration while I constructed the experiment presented in this thesis. Michael was always encouraging, even when progress was slow, and he never wavered in his conviction that everything would come together in the end.

I would like to thank my defense committee members, Prof. Alireza Marandi, Prof. Jason Alicea, and Prof. Jim Crutchfield for their attention and feedback on this work. I would also like to thank Prof. Chiara Daraio for her feedback on my candidacy exam.

This work would not have been possible without my sync co-conspirators who helped and guided me over the last six years. I would like to thank Dr. Damien Maillard, Prof. Guillermo Villanueva, Dr. Warren Fon, Dr. Matthew Matheny, Prof. Jim Crutchfield, and Isaiah Curtis. I am grateful to Damien and Guillermo for their meticulous efforts towards fabricating the piezoelectric devices used in this work. It was a great pleasure to work with them, and, without their assistance, this project may have never come to fruition. I am grateful to Matt and Warren for allowing me to take over their project and for sharing all of their synchronization wisdom with me, even after departing from Caltech. I am thankful to Jim for his endless enthusiasm and patience as I worked out many experimental kinks and for bringing me to multiple wonderful Telluride workshops where I met and engaged with many great scientists. Finally, many thanks to my friend and colleague, Isaiah, for joining the team and contributing many insightful ideas about analysis of and computing with NEMS networks. It is my honor to have worked with such a talented group of people.

I would also like to extend my gratitude to the entire NaPSAC team — Prof. Philip Feng, Prof. Ken Sheppard, Prof. Sarah Marzen, Prof. Jing Guo, Tim Caplice, Dr. Taeju Lee, Dr. Tahmid Kaiser, Guoting, Sariha, and Byoung-Soo. It was a great pleasure working with you on the next generation of integrated NEMS networks.

It has also been my pleasure to meet many great scientists during my time at Caltech, and I would like to extend my gratitude for many fruitful conversations that I have had with Prof. Steve Shaw, Prof. John Sader, Dr. Alec Boyd, Dr. Kyle Ray, Chris Pratt, and Dr. Greg Wimsatt.

From my time before Caltech, I would like to thank Dr. Ben Griffin, Prof. Jeff Rhoads, and Prof. Nikhil Bajaj for mentoring me early in my career and for inspiring me to continue with research.

I would like to thank all of the Roukes group (past and present) for many great discussions, help around the lab, and for tolerating my persistent heckling: Alex, Nitay, Teresa, Saif, Cenk, Mark, Dvin, Atakan, Xinchang, Arda, Batu, Andrew, Ewa, Raj, Marion, Jeff, Lo, Niki, Alice, Sophie, Tom, Alfredo, and Kirill. I would also like to thank to all of the Caltech staff who made the journey as smooth as possible: Tracy, Trudy, Stephanie, Loly, Jennifer, Christy, Guy, Kelly, Alex, Bert, Tiffany, Lena, and Nancy, among numerous others.

I feel immensely fortunate to have met so many wonderful people from all around the world during my time at Caltech. I am constantly inspired by who my friends are and proud of them for all they have accomplished. My time in graduate school was far richer because of them. Thank you to Mert, Utku, Hirsh, Dulce, Tanner, Jess, Brayden, Kate, Aki, Nahoko, Sahangi, Shruti, Yanky, André, Andrea, Sari, Prachi, Flavia, Andrea, Patty, Andrea, Alkim, Nikos, Ramon, Fredrik, Preston, Guglielmo, Filip, Olli, Jean, Dima, Efe, Kailah, Marianne, Ina, Liam, Robby, Komron, Morgan, Jag, John, Andreas, Piero, Steven, Samir, Sameer, Dhruv, Nick, Robby, Erik, Taylan, Till, Brecht, Cat, Kaustubh, Garima, Bilge, Haakon, Jonas, Alex, Arun, Branislav, and Jonah. And to my many friends outside of Caltech who have always been there to share a laugh: Olivia, Cam, Tom, Rachel, Ethan, Jon, Andromeda, Nic, Mark, Madi, Randon, Daniel, JC, Katie, Christine, Brian, Courtney, Joe, Nicole, and Brandon.

I thank all of my extended family, grandparents, aunts, uncles, and cousins, for their constant support and love, and for always putting up with my inability to sit still.

And, finally, I would like to thank my mom, dad, and brother. You have always been there for me and supported me, which enabled me to pursue and complete this PhD. I am inspired by the passion that you put into your own work and thank you for your steadfast love and support during this journey.

ABSTRACT

Networks of nonlinear limit-cycle oscillators exhibit a broad range of collective dynamical behaviors, from simple synchronization to exotic, partially phase-locked chimera states. Such dynamics are universal to many natural and technological systems and may prove useful as building blocks for special-purpose analog computers. In this thesis, we present the design, construction, and analysis of a network of eight nonlinear nanoelectromechanical oscillators with electronic feedback and coupling. Precise control over the system parameters, including the nodal frequencies, nodal nonlinearities, inter-nodal coupling strengths, and network topology, enables the systematic study of synchronization, pattern formation, and computation in complex networks with real-world dynamics, parameter disorder, and noise. While investigating the network's dynamics on different graph structures — namely, all 8-node circulant graphs, a wheel graph, and Erdős–Rényi random graphs — we observed a rich set of synchronized and unsynchronized states, arising even in the weak coupling limit. Comparison with an accompanying theoretical model reveals mechanisms that cause physical networks to deviate from idealized behavior. Finally, we report recent efforts to scale the network to a larger number of nodes by integrating an array of nanoelectromechanical resonators with an application-specific integrated circuit.

TABLE OF CONTENTS

Acknowledgements iv
Abstract vi
Table of Contents vii
List of Illustrations ix
List of Tables xi
Chapter I: Introduction 1
Chapter II: Dynamics of Coupled Oscillators 3
2.1 Basic Properties of Oscillators 3
2.2 Canonical Models 6
2.3 Reactively Coupled Oscillators with Saturated Feedback 10
2.4 Parametrically Driven Nonlinear Oscillator 13
2.5 Previous Experiments: Exotic States on a Ring 13
2.6 Conclusion 14
Chapter III: Computing with Nanomechanical Oscillators 15
3.1 Background 15
3.2 Annealing for Combinatorial Optimization 17
3.3 Reservoir Computing 20
3.4 Conclusion 26
Chapter IV: Piezoelectric, Flexural-Mode NEMS Resonators 27
4.1 Background 27
4.2 Fabrication 40
4.3 Single Device Characterization 41
4.4 Ensemble Characterization 46
4.5 Conclusion 48
Chapter V: Oscillators with Electronic Feedback and Coupling 49
5.1 Network Hardware Overview 50
5.2 Nonlinear Oscillators with Multiple Feedback Loops 54
5.3 Experimental Control Parameters 62
5.4 Node Calibration (Experiment) 63
5.5 Edge Calibration (Experiment) 68
5.6 Conclusion 73
Chapter VI: Synchronized and Unsynchronized States of Uniform Complex Graphs 74
6.1 Experiment Protocol 74
6.2 Circulant Graphs 75
6.3 Wheel Graph 82
6.4 Erdős–Rényi Graphs 84
6.5 Conclusion 87
Chapter VII: Towards NEMS Networks with ASIC Technology 88

7.1 GHz NEMS Arrays . . . 88
7.2 Flip-Chip Integration . . . 90
7.3 Conclusion . . . 91
Chapter VIII: Conclusion . . . 93
Appendix A: Derivation of Slow-Time Equations of Motion for Weakly Coupled Oscillators . . . 95
A.1 Solution via Secular Perturbation Theory . . . 97
Appendix B: Derivation of Higher Order Phase Dynamics in the Small Coupling Limit . . . 104
B.1 Intro . . . 104
B.2 Dissipative Coupling . . . 105
B.3 Reactive Coupling . . . 109
Appendix C: Dynamics of Two Coupled Oscillators . . . 112
C.1 Dissipative Coupling . . . 112
C.2 Reactive Coupling . . . 114
C.3 Reactive vs. Dissipative Coupling . . . 115
Appendix D: Additional Device Characterization . . . 117
D.1 Large Amplitude Nonlinear Response . . . 117
D.2 Frequency Response vs DC Bias . . . 118
D.3 Die and Wafer Level Statistics . . . 119
Appendix E: Detailed Experimental Setup . . . 123
E.1 Hardware Overview . . . 123
E.2 Network Control . . . 125
E.3 Experimental Network Parameters . . . 126
Appendix F: Additional Experimental Methods . . . 127
F.1 The Hilbert Transform . . . 127
F.2 Mutual Information . . . 128
F.3 Calculating Allan Deviation . . . 129

LIST OF ILLUSTRATIONS

Number *Page*

2.1 **Van der Pol Oscillator Flow** . . . 4
2.2 **Explosive Synchronization** . . . 9
2.3 **Higher Order Interactions** . . . 12
3.1 **Reservoir Computing Overview** . . . 20
3.2 **Explosive Kuramoto RC λ Sweep** . . . 25
4.1 **Device SEM and Material Stack** . . . 27
4.2 **Nonlinear Potential** . . . 31
4.3 **Theoretical Duffing Response** . . . 34
4.4 **Piezoelectric Actuation and Detection** . . . 36
4.5 **mBVD Model** . . . 38
4.6 **Responsivity Model** . . . 38
4.7 **Probe Station Measurement Setup** . . . 42
4.8 **Single Device Nonlinear Frequency Response** . . . 43
4.9 **Backbone Fit** . . . 44
4.10 **DC Tuning Frequency Response** . . . 45
4.11 **Device Properties vs DC Bias** . . . 46
4.12 **Ensemble Device Characterization** . . . 48
5.1 **Two Oscillator Coupling Schematic** . . . 51
5.2 **Eight Oscillator Coupling Schematic** . . . 53
5.3 **Response to Feedback Phase (Theory)** . . . 56
5.4 **Linear Response to Coupling Phase (Theory)** . . . 59
5.5 **Nonlinear Response to Coupling Phase (Theory)** . . . 60
5.6 **Types of Linear Coupling** . . . 61
5.7 **Feedback Phase Calibration (Experiment)** . . . 64
5.8 **Oscillator Nonlinearity Calibration (Experiment)** . . . 66
5.9 **Oscillator Characteristics** . . . 66
5.10 **Wide PSD** . . . 67
5.11 **Amplitude and Frequency Fluctuations** . . . 69
5.12 **Parametric Sweep of Edge Phase (Experiment)** . . . 71
5.13 **$|\beta|$ Calibration (Experiment)** . . . 72
6.1 **Time Domain Synchronization on a Ring** . . . 75

6.2 **8-Node Circulant Graphs** . . . 76
6.3 **Synchronized States of a Ring (Group 1)** . . . 77
6.4 **States of Second-Order Circulant Graphs (Group 2)** . . . 78
6.5 $C_8(\{1,2\})$ **Decoupled State Details** . . . 79
6.6 **States of Third-Order Circulant Graphs (Group 3)** . . . 80
6.7 C_{123} **State Details** . . . 81
6.8 **Complete Graph States (Group 4)** . . . 82
6.9 **Wheel Graph** . . . 82
6.10 **8-Node Wheel Graph States** . . . 83
6.11 **Synchronized States of an ER Graph with Reactive-Repulsive Coupling** . . . 85
6.12 **Synchronized States of an ER Graph with Reactive-Attractive Coupling** . . . 86
7.1 **ASIC on NEMS Vertical Integration Design** . . . 89
7.2 **NEMS Array** . . . 90
7.3 **NEMS Response and Frequency Histogram** . . . 91
7.4 **ASIC on NEMS Vertical Integration Test** . . . 92
7.5 **Pictures of Mechanical Integration** . . . 92
C.1 **Locking Dynamics of Two Coupled Oscillators** . . . 116
D.1 **Resonator Response at Large Drive Power (Linear Amplitude)** . . . 118
D.2 **Resonator Response at Large Drive Power (Logarithmic)** . . . 119
D.3 **Resonator Transmission (S21)** . . . 120
D.4 **Raw Data of DC and Amplitude Sweep** . . . 121
D.5 **Measured Backbone Curves for each DC Bias** . . . 121
D.6 **Resonant Frequency vs Device Size** . . . 122
D.7 **Frequency and Q Histograms** . . . 122
E.1 **Experiment Hardware** . . . 124
E.2 **Oscillator PCB** . . . 124
E.3 **Control and Readout** . . . 125

LIST OF TABLES

Number *Page*

4.1 Summary of Device Parameters 47
5.1 Equations for $\bar{a}$ and Ω at Critical Values of Δ 62
6.1 Counts of State (c) vs. State (d) for Varying β (Experimental Data). . 85
A.1 Exact vs. Approximate ω_n at Selected Detunings and with Illustrative Parameter Values ($f_0 = 1.8\,\text{MHz}$, $Q = 2500$, $\Delta f_{3\text{dB}} = 720\,\text{Hz}$). 97
C.1 Summary of Dynamics of Two Coupled Oscillators 115
E.1 Properties of the NEMS Resonators 126

Chapter 1

INTRODUCTION

The study of networks is ubiquitous in modern science and engineering. They are examined in various contexts, ranging from biological networks including the human brain or synchronous fireflies in Southeast Asia, to engineered systems, such as the power grid and artificial neural networks. Many of these high-dimensional and often nonlinear systems are self-oscillatory, where each node in the network is an independent oscillator which periodically varies in time according to its intrinsic rhythm. Individually, each node is rather boring, and the basic properties of individual oscillators have long been understood. However, when interconnected in a network, coupled oscillators can exhibit fascinating phenomena which can vary widely depending on the parameters and conditions of the network. On one hand, they can synchronize, where a network of individual oscillators with disordered intrinsic frequencies and weak coupling can reach a global consensus and oscillate in lockstep. On the other hand, strongly coupled oscillators which are intrinsically identical can break into complex symmetry-broken patterns where they evolve in clustered and partially synchronized states or even support nonlinear wave propagation.

Understanding these simple and exotic behaviors of oscillatory networks is a paramount scientific challenge, and progress in the field promises a myriad of applications including a better understanding of the human brain and heart, a more efficient or robust power grid, and low-power analog computers for solving complex optimization or machine learning problems. Most studies of real-world networks done to date are top-down, where we try to understand the behavior of existing networks by measuring and modeling them. Often, this comes with several challenges, some of which include the lack of full state readout, limited knowledge about the connectivity, very slow or very fast dynamical timescales (sometimes both), and limited control over the nodes or edges in the network. As a result, much of the research on coupled oscillators done today is computational; however, while intrinsically interesting, there is often little to no experimental data to compare these explorations to or to guide them. Current experimental platforms lack the necessary rigor for a deep exploration of the parameter space, so bottom-up experimental paradigms with

excellent control over the parameters and which operate under real-world conditions must be constructed.

In addition to the fundamental science, experiments are increasingly important for realizing practical applications, such as oscillator-based computing. Under this paradigm, the proposal is that by using analog oscillators to do computation, we can achieve faster compute times, lower power, lower footprint, increased robustness, or increased security, among other potential advantages. Similar to quantum computing, there are a number of hardware platforms which have been proposed for building utility-scale oscillatory computers, including photonics, nanomechanics, optomechanics, spintronics, and CMOS electronics, to name a few. It remains unclear which, if any, of these technologies will prove useful across the wide range of application spaces, although some are more advanced and promising than others.

Our goal in this thesis is to study the fundamental dynamics of real-world nonlinear oscillator networks with an eye towards potential applications in computing. Towards that goal, we built an analog network of 8 NEMS-based oscillators with electronically tunable feedback and coupling. Leveraging our nanomechanical resonators, each node is frequency tunable by more than 50 linewidths and has a nodal nonlinearity tunable by more than 32 dB. The topology is fully controllable on µs timescales and can realize any adjacency matrix. For a subset of the edges, the strength can be tuned over a dynamic range of more than 47 dB and the phase can be controlled over a 160° window. In this thesis, we cover the design and construction of the network and investigate its dynamics with various complex topologies and over a range of nodal and inter-nodal parameters.

In Chapter 2 we review the basic dynamics of some widely studied models of coupled oscillators. Then, in Chapter 3 we explore some proposals for computing with coupled oscillators, including reservoir computing and combinatorial optimization. Chapter 4 and Chapter 5 cover the NEMS devices and network architecture, respectively. Then, in Chapter 6, we report and analyze experimentally observed synchronized states of complex graphs, including all circulant graphs of 8 nodes, a wheel graph, and various random graphs. In Chapter 7 we detail recent efforts to integrate a NEMS array with an ASIC to achieve large-scale and low-power coupled oscillator networks at GHz frequencies. Finally, in Chapter 8, we conclude this work with some remarks about NEMS networks and analog oscillator computation.

Chapter 2

DYNAMICS OF COUPLED OSCILLATORS

One fascinating feature of coupled oscillators is that the dynamics can range from the remarkably simple to the remarkably complex. On one hand, networks with parameter disorder spontaneously synchronize into phase-locked states whereby the entire network collapses to and evolves along a one-dimensional manifold. On the other hand, an initially homogeneous network can spontaneously break its underlying spatial symmetries, allowing it to evolve on high-dimensional manifolds where exotic states abound. Beyond steady-state operation, the transient dynamics of networks can also be highly complex and span multiple timescales.

It would be nearly impossible to cover the vast range of coupled oscillator models or to examine all of the possible dynamics. So, in this chapter, we aim to introduce the dynamics of coupled oscillators in a historical and phenomenological way with an eye towards the experiments that will be covered in this thesis. First, we introduce what an oscillator is and what its key properties are. Then, we cover dynamics exhibited by the most commonly used amplitude-phase and phase-only models. We specifically introduce two niche phenomena, the explosive synchronization transition and the parametrically driven feedback oscillator.

2.1 Basic Properties of Oscillators

Externally driven resonant systems and oscillators are often referred to as driven-dissipative systems and gain-dissipative systems, respectively. This distinction is important because they exhibit different classes of behavior, which are often conflated in the literature. In this work, we principally study the behavior of limit cycle oscillators — henceforth referred to as oscillators —, a common class of gain-dissipative system. Most important for this work is the property that oscillators self-adjust their frequency under the influence of weak interactions while driven resonators do not. To build intuition about oscillatory systems, we introduce the canonical van der Pol oscillator [1]:

$$\ddot{x} + \epsilon(x^2 - 1)\dot{x} + x = 0. \tag{2.1}$$

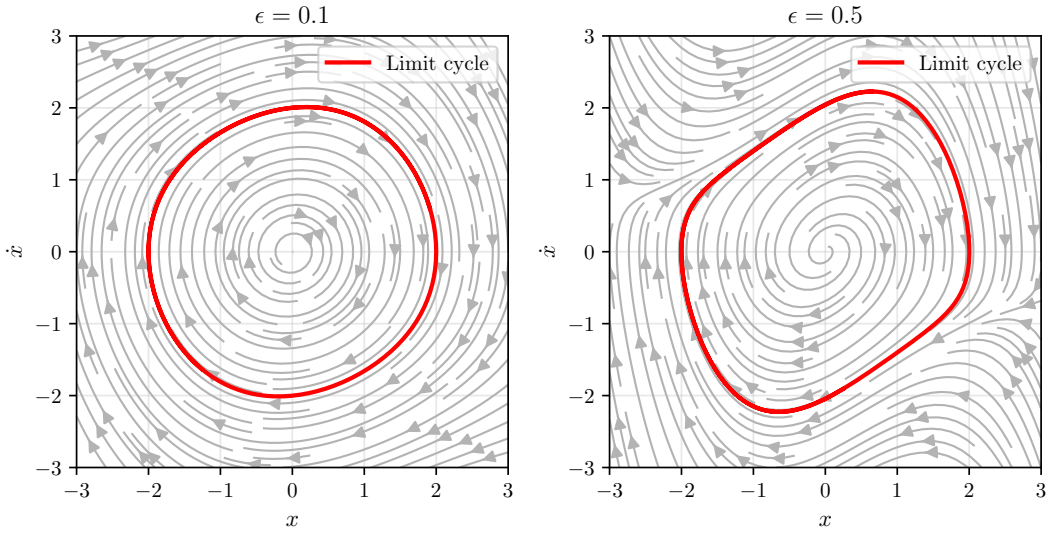

Figure 2.1: Van der Pol Oscillator Flow. 2D flow map of the van der Pol oscillator for (left) $\epsilon = 0.1$ and (right) $\epsilon = 0.5$. Red lines indicate the limit cycle for each case.

Here, x is the position coordinate of an arbitrary system. This equation describes an undriven second-order system which has a nonlinear damping, described by $\epsilon(x^2 - 1)$. Rewriting the system in state space coordinates, x_1 and x_2 gives

$$x_1 = x \tag{2.2}$$

$$x_2 = \dot{x} \tag{2.3}$$

$$\dot{x}_1 = \dot{x} = x_2 \tag{2.4}$$

$$\dot{x}_2 = \ddot{x} = -\epsilon(x_1^2 - 1)x_2 - x_1. \tag{2.5}$$

Using these equations, we plot the 2D flow map in Fig. 2.1 for $\epsilon = 0.1$ and $\epsilon = 0.5$. These plots show the streamlines for the system, which indicate the direction the system will travel given a specific x and $\dot{x}$. The speed of the flow is not indicated. On each plot, the red trajectory denotes the limit cycle, which is a stable attractor of the system. Given any initial condition, except for the unstable fixed point of $x = \dot{x} = 0$, the system will converge to the limit cycle. Trajectories starting inside grow outward to it while trajectories starting outside decay inward to it. This demonstrates the stability of the limit cycle's amplitude, which can be characterized by the local Lyapunov exponents. The amount of nonlinear damping changes the rate at which trajectories converge to the limit cycle and the shape, but the limit cycle always persists for $\epsilon > 0$.

The van der Pol oscillator (2.1) illustrates a number of key properties applicable to many classes of oscillatory systems. Firstly, it has a nonlinear damping. All oscillators require a nonlinearity to limit the amplitude. For direct feedback oscillators, this typically arises from nonlinear feedback but can also arise from intrinsic nonlinear damping of the resonator. For relaxation oscillators, this appears in the form of a nonlinear switch which is activated upon reaching a certain threshold and exhibits some hysteresis.

Furthermore, it is an autonomous dynamical system, meaning that it is invariant under time translation. Because phase is a function of time, the system is also invariant under phase translation resulting in a property known as phase freedom. In other words, the system is neutrally stable in phase, the direction tangential to the limit cycle and characterized by a zero Lyapunov exponent. Consequently, perturbations to the phase persist in time. This allows oscillators to adjust their frequency under the weak influence of other oscillators or external injection signals, allowing them to synchronize. This behavior is in contrast to driven resonators which have a preferred value of phase dictated by the drive signal and the system's phase response. This precludes driven resonators from synchronization as it is traditionally defined in the field of nonlinear dynamics [2].

> Oscillators have phase freedom which allows them to adjust their frequencies and synchronize under weak coupling to other oscillators or periodic injection signals. This is in contrast to driven resonators, which have a preferred value of phase dictated by the drive and the resonator's phase response.

Beyond limit cycle oscillators and driven resonators, there are some edge cases which are more challenging to classify, but we briefly address the most common ones. Firstly, the simple harmonic oscillator (SHO), given by (2.6):

$$\ddot{x} + x = 0. \tag{2.6}$$

The SHO is a resonator with no dissipation. As a result, it can oscillate forever given some initial energy. It is an autonomous differential equation, but it also has a neutrally stable amplitude. Most relevantly, it is a conservative system and under coherent coupling, SHOs cannot synchronize. Instead, a system of N coupled SHOs generates a set of N non-dissipative normal modes.

The second fringe case is coupled resonators which have some damping but are not driven, such as reported in [3]. Here, the resonators are linear and damped, and, given some initial energy, their phases can attract under dissipative coupling. However, the nodes are not self-sustaining. Therefore, it is not clear how to classify this phenomenon based on the widely held definition of synchronization.

2.2 Canonical Models

Here we introduce a few mathematical models of coupled oscillators which will be referenced throughout this thesis.

Stuart–Landau Model

All real-world oscillators have both an amplitude and a phase. Certain classes of oscillators can be immediately reduced to phase-only description, such as CMOS ring oscillators, which have highly saturated amplitudes effectively pinning them at a specific voltage. However, a large class of oscillators have amplitude and phase dynamics and only under certain conditions can be accurately modeled by a phase-only description. Many oscillators can be directly described by the canonical Stuart–Landau oscillator model, which formulates the amplitude and phase of a feedback oscillator operating near a Hopf bifurcation:

$$\frac{dA}{dT} = (\tilde{\gamma} + i\omega)A - \frac{1}{2}(\eta + i\alpha)A|A|^2 \tag{2.7}$$

$$= \left(\tilde{\gamma} - \eta|A|^2\right) A + i\left(\omega - \alpha|A|^2\right) A \tag{2.8}$$

where A is a complex-valued field, $\tilde{\gamma}$ is the linear dissipation, η is the quadratic nonlinear dissipation, ω is the linear frequency, and α is the quadratic amplitude frequency coupling. Networks can then be constructed by introducing linear or nonlinear inter-oscillator couplings. Such networks have been widely studied and can produce a wide range of dynamics ranging from simple synchronization to decoupling [4], explosive oscillation death [5], and periodic waves under uni-directional coupling with time delay [6]. In the weak coupling limit, the Stuart–Landau model can be understood through phase reduction techniques [7].

The Kuramoto Model

In 1975, the mathematician Yoshiki Kuramoto introduced the now foundational model of synchronization [8]. His insight was to assume that the linear dissipation

and nonlinear feedback from the Stuart–Landau oscillator model were infinitely large but equal, allowing the nodal amplitude dynamics to relax infinitely fast. With the addition of a dissipative mean-field coupling, he derived a first-order phase model which had only a sine coupling function:

$$\frac{d\phi_n}{dt} = \omega_n + \frac{K}{N}\sum_{m=1}^{N}\sin(\phi_m - \phi_n) \tag{2.9}$$

where ϕ_n is the phase of oscillator n, ω_n is the normalized frequency of oscillator n, K is the global coupling strength, and N is the number of oscillators. By considering an infinite population of mean-field coupled oscillators, and defining the order parameter, r, such that

$$re^{i\psi} = \frac{1}{N}\sum_{n=1}^{N} e^{i\phi_n}, \tag{2.10}$$

Kuramoto was able to exactly solve for r as a function of the coupling strength, K, for symmetric distributions of ω_n [8, 9]. What he discovered is that there is a second-order phase transition from an incoherent, unsynchronized state ($r = 0$) to a frequency synchronized state ($r > 0$) at a critical value of the coupling strength. Above the critical value, more and more oscillators become frequency entrained with the coherent group until they all reach a common frequency [10]. This critical value of the onset of synchronization is a function of the spread of the frequencies.

The Generalized Kuramoto Model

The Kuramoto model (2.9) can be generalized to arbitrary networks and written as

$$\frac{d\phi_n}{dt} = \omega_n + \sum_{m=1}^{N} K_{nm}\sin(\phi_m - \phi_n) \tag{2.11}$$

where $K_{nm} \in \mathbb{R}^{N\times N}$ defines an arbitrary adjacency matrix and coupling strengths, relaxing the conditions of uniform mean-field coupling. Both (2.9) and (2.11) have been extensively studied for different network topologies, variations of the coupling strengths, and variations of the frequency distribution.

One simplification that has been explored is the case of identical natural frequencies: $\omega_n = \omega \,\forall\, n$. Noting that the average natural frequency can always be removed by applying the transformation $\phi_n \to \phi_n - \bar{\omega}t$, where $\bar{\omega}$ is the mean natural frequency.

This case is unique because the dynamics are a gradient system [11]; thus, we can define a Lyapunov energy function of the form

$$V = -\sum_{m=1}^{N} K_{nm} \cos(\phi_m - \phi_n) \tag{2.12}$$

such that $d\phi/dt = -\partial V/\partial \phi_i$. Importantly, as gradient systems evolve, they always descend in this potential landscape; therefore, only fixed point solutions exist. Even with this level of simplification, there are still unanswered mathematical questions. For example, there is a long line of research searching for the limit of connection density at which only the in-phase synchronized state is stable [12–16]. The gradient nature of this dynamic also makes it amenable for different oscillator-based computing schemes which rely on energy minimization, as will be introduced in Chapter 3. For a more detailed review of the history and developments around the Kuramoto model, see [17].

Explosive Kuramoto Model

By ensuring the proper parameter settings and network conditions, it is possible to engineer first-order, explosive phase transitions in coupled oscillator networks [18–22]. First-order transitions are characterized by an abrupt jump from disorder to order at a critical coupling strength. Typically, such a transition is manifested by weighting the coupling strength to a node according to the nodes connectivity, natural frequency, or both. Explosive phenomena are interesting because they manifest as extreme state changes resulting from small, continuous parameter variation. Such criticality may play an important role in many natural phenomena, such as the spread of an epidemic [23], or may also be useful for engineering critical dynamics for computation [24]. There are many possible constructions, but one interesting model which was studied in [20] is given by

$$\frac{d\phi_n}{dt} = \omega_n + \frac{\lambda|\omega_n|}{k_n} \sum_{m=1}^{N} K_{nm} \sin(\phi_m - \phi_n) \tag{2.13}$$

where λ is a global coupling strength and k_n is the degree of node n. Here, the coupling is weighted by each oscillators intrinsic frequency, ω_n. By choosing ω from a normal distribution with standard deviation σ, such that $\omega \sim \mathcal{N}(0, \sigma^2)$ and K_{nm} to be an Erdős–Rényi random graph, this model can support explosive synchronization. An Erdős–Rényi graph is a random graph where each edge between N nodes has an

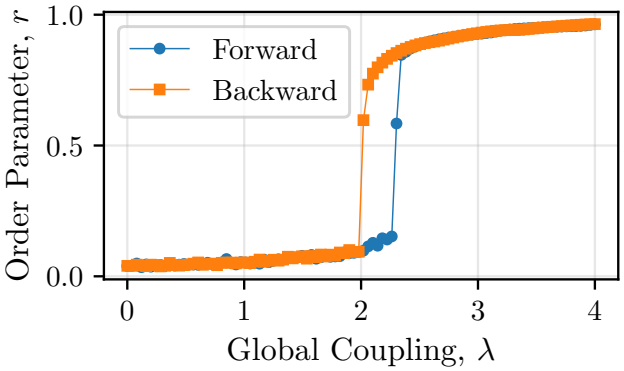

Figure 2.2: Explosive Synchronization. Up and down sweep of the global coupling parameter demonstrating explosive synchronization.

independent probability of being created, given by p. To illustrate the dynamics, we generate an Erdős–Rényi graph with 500 nodes and $p = 0.2$, giving an average node degree of approximately $k_n = 40$. Using this graph, we simulate the dynamics with $\sigma = 1$ while increasing and then decreasing λ between 0 and 4. Plotting the resulting order parameter, r, in Fig. 2.2, we observe a clear region of hysteresis, indicating a first-order transition.

This is just one example of an explosive transition and there are many other possible models exhibiting such dynamics [25]. In this case, the first-order transition is robust and exists across a wide range of the parameter space.

The Generalized Kuramoto–Sakaguchi Model

In 1986, Kuramoto and Sakaguchi extended their model to include a constant phase shift in the coupling [26]. The generalized version, which we will refer to as the Kuramoto–Sakaguchi (KS) model is given by (2.14)

$$\frac{d\phi_n}{dt} = \omega_n + \sum_{m=1}^{N} K_{nm} \sin(\phi_m - \phi_n + \psi) \tag{2.14}$$

$$= \omega_n + \sum_{m=1}^{N} \tilde{K}_{nm} \Big[\sin(\phi_m - \phi_n) + \gamma \cos(\phi_m - \phi_n) \Big], \tag{2.15}$$

where ψ is the phase shift of the coupling. The equation can be rewritten as (2.15) where $\tilde{K}_{nm} = K_{nm} \cos\psi$ and $\gamma = \tan\psi$. This model is no longer a gradient system due to the cosine coupling term. Thus, even with uniform ω_n, its solutions can include more complex, nonstationary attractors, as opposed to only fixed points. This

phase lag can arise from real-time delay in the interaction between two oscillators, or it can manifest naturally in amplitude-phase models which have no time delay. We will see how this arises from phase reduction in the next sections.

2.3 Reactively Coupled Oscillators with Saturated Feedback

Finally, we introduce a model for a saturated feedback oscillator, which will describe the experimental system used in this work. It has been previously studied in [27, 28] and is similar to the Stuart–Landau model but with a saturated feedback instead of a quadratic nonlinear dissipation. The model is derived in Appendix A by using secular perturbation theory:

$$\frac{dA_n}{dT} = \frac{1}{2}\frac{A_n}{|A_n|} - \frac{1}{2}A_n + i\left(\omega_n + \alpha_n|A_n|^2\right)A_n + \frac{\beta}{2}\sum_{m=1}^{N}\frac{C_{nm}}{k_n}A_m \tag{2.16}$$

where $A_n = a_n e^{i\phi_n}$ is the complex value representing the phase and amplitude of oscillator n, ω_n is the normalized intrinsic frequency deviation of node n (see Appendix A for details), α_n is the quadratic amplitude-frequency coupling strength, $\beta \in \mathbb{C}$ is the inter-node coupling strength, C_{nm} is the inter-node adjacency matrix, and $k_n = \sum_m^N C_{nm}$ is the input degree of node n. We examine the case of purely reactive coupling, which corresponds to β being purely imaginary by the convention adopted in (2.16). Additionally, we let $\lambda = -\mathrm{Im}(\beta)$ be the magnitude of the reactive coupling strength so that positive λ defines attractive coupling (see Appendix B). After splitting the equation into real and imaginary components, we get two real-valued, coupled equations for the amplitude and phase

$$a_n' = \frac{1}{2} - \frac{1}{2}a_n + \frac{\lambda}{2}\sum_{m=1}^{N}\frac{C_{nm}}{k_n}a_m\sin(\phi_m - \phi_n) \tag{2.17}$$

$$\phi_n' = \omega_n + \alpha_n a_n^2 - \frac{\lambda}{2}\sum_{m=1}^{N}\frac{C_{nm}}{k_n}\frac{a_m}{a_n}\cos(\phi_m - \phi_n). \tag{2.18}$$

Higher-Order Phase Reduction

To elucidate some of the dynamics, we perform a phase reduction by adiabatic elimination of the fast amplitude variations. The scaled relaxation rate is equal to $1/2$ in (2.17). In the limit of weak coupling ($\lambda << 1$), we can write the amplitude and phase of the oscillators as a power series in the coupling strength, λ:

$$a_n = a_n^{(0)} + \lambda a_n^{(1)} + \lambda^2 a_n^{(2)} + O(\lambda^3) \tag{2.19}$$

$$\phi_n = \phi_n^{(0)} + \lambda \phi_n^{(1)} + \lambda^2 \phi_n^{(2)} + O(\lambda^3). \tag{2.20}$$

Plugging (2.19) and (2.20) into (2.17) and enforcing the fast magnitude relaxation, we derive a second-order correction to the amplitude (2.21) which describes slow amplitude variations resulting from the phase dynamics:

$$a_n = 1 + \lambda \sum_m^N \frac{C_{nm}}{k_n} \sin(\phi_m - \phi_n) + \lambda^2 \sum_{m,k}^N \frac{C_{nm}}{k_n} \frac{C_{mk}}{k_m} \sin(\phi_k - \phi_m) \sin(\phi_m - \phi_n) + O(\lambda^3). \tag{2.21}$$

Plugging (2.20) and (2.21) into (2.18) yields the higher-order phase dynamics (2.22) up to $O(\lambda^2)$:

$$\begin{aligned}
\phi_n' &= \omega_n + \alpha_n \\
&+ \frac{\lambda}{2} \sum_m^N \frac{C_{nm}}{k_n} \Big\{ 4\alpha_n \sin(\phi_m - \phi_n) - \cos(\phi_m - \phi_n) \Big\} \\
&+ \frac{\lambda^2}{4} \sum_m \frac{C_{nm}^2}{k_n^2} \Big\{ 2\alpha_n - 2\alpha_n \cos\left[2(\phi_m - \phi_n)\right] + \sin[2(\phi_m - \phi_n)] \Big\} \\
&+ \frac{\lambda^2}{4} \sum_m \sum_{k \neq m} \frac{C_{nm} C_{nk}}{k_n^2} \Big\{ 2\alpha_n \cos(\phi_k - \phi_m) - 2\alpha_n \cos(\phi_k + \phi_m - 2\phi_n) \\
&\qquad + \sin(\phi_k + \phi_m - 2\phi_n) + \sin(\phi_k - \phi_m) \Big\} \\
&+ \frac{\lambda^2}{4} \sum_m \sum_{k \neq m} \frac{C_{nm}}{k_n} \frac{C_{mk}}{k_m} \Big\{ - 4\alpha_n \cos(\phi_k - \phi_n) + 4\alpha_n \cos(\phi_k + \phi_n - 2\phi_m) \\
&\qquad - \sin(\phi_k + \phi_n - 2\phi_m) - \sin(\phi_k - \phi_n) \Big\}.
\end{aligned} \tag{2.22}$$

This expansion reduces the system to a higher-order phase model valid for any adjacency matrix, C. The full derivation is included in Appendix B. By truncating at the first order, we get the Kuramoto–Sakaguchi coupling, with dynamics given by

$$\phi_n' = \omega_n + \alpha_n + \frac{\lambda}{2} \sum_m^N \frac{C_{nm}}{k_n} \Big\{ 4\alpha_n \sin(\phi_m - \phi_n) - \cos(\phi_m - \phi_n) \Big\}. \tag{2.23}$$

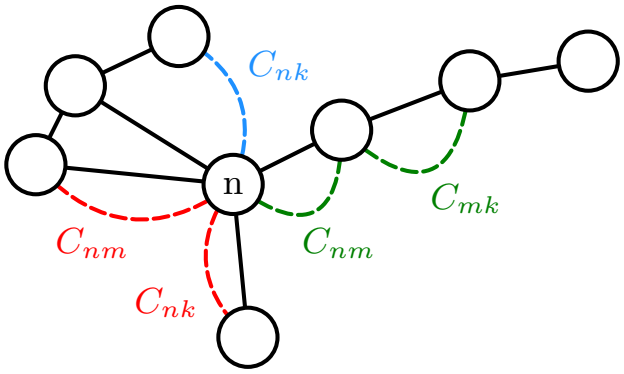

Figure 2.3: Higher Order Interactions. Schematic depiction of higher-order interactions in the phase dynamics.

From the equation, we see that α_n induces dissipative phase dynamics, and, for large α_n the system reduces to the Kuramoto model. At higher orders, there are emergent multibody interactions between the oscillators, appearing as next-nearest-neighbor, biharmonic, and triadic terms, which are illustrated in Fig. 2.3. Some terms arise as a result of nodal nonlinearity while the others arise solely from the inter-nodal coupling. The equation contains nearest-neighbor (NN), biharmonic NN, next-nearest-neighbor (NNN), triadic centered, triadic offset, and 'other node' terms.

The form of each term can be understood by examining the coefficient. The fourth line has a coefficient of $C_{nm}C_{nk}$ which describes triadic terms containing node n and two of its nearest neighbors, m and k, if those connections are present. Interestingly, there is also a term which only contains the phases of nodes k and m. The last line has a coefficient $C_{nm}C_{mk}$ which describes triadic terms where there is a connection between n and m plus a connection between m and another node, k, if those connections are present in the topology. Here, we also see direct NNN couplings between n and k.

As was shown in [28], these higher-order terms can stabilize exotic states in simple networks, such as a ring. For complex networks, we expect that as the coupling strength increases, the higher-order terms should play a role in the network dynamics.

2.4 Parametrically Driven Nonlinear Oscillator

So far, we have only introduced an autonomous model for coupled oscillators. But, in addition to self-feedback, an oscillator can simultaneously be parametrically driven by external signals. If those external signals are at integer multiples of the oscillation frequency, the dynamics can be described by [29]

$$\frac{dA}{dT} = \frac{1}{2}\frac{A}{|A|} - \frac{1}{2}A + i\alpha|A|^2A + h_nA^{*(n-1)}e^{i\Omega_nT} \tag{2.24}$$

where A^* is the complex conjugate of A, and Ω_n is detuning of the drive from the n^{th} harmonic. In the case of $n = 1$, the oscillator is subjected to a resonant drive which injection locks the oscillator. This gives the system a preferred orientation of the phase. For $n = 2, 3, 4...$, the conjugate term describes subharmonic injection locking of the oscillator. Under this forcing, the oscillator can have multiple preferred phase states depending on the order of the injection locking.

2.5 Previous Experiments: Exotic States on a Ring

The specific case of a ring was thoroughly studied in experiment by Matheny et al. [28] and has been the subject of many theoretical papers, such as [4]. In [28], the authors explored the dynamics across a wide range of nodal nonlinearities, α, and coupling strengths, $|\beta|$, with reactive-repulsive coupling. They found a variety of experimentally observable exotic states including anti-phase and in-phase synchronization, splay states (rotating waves), inhomogeneous synchronization, decoupled states, weak chimeras, and traveling waves, among other more complex phenomena such as a noise driven chimera. Using group theoretical analysis, the authors showed that the states arise as a result of spontaneous symmetry breaking and each state, with a few exceptions, is associated with an isotropy subgroup of the underlying physical couplings — a ring — composed with time translation symmetry (the 1-torus). In other words, each state is invariant under the action of the generator of the subgroup associated with it, defining the invariant subspace. A key finding of the paper is that not all of these states appear to be stable for large α and small $|\beta|$; instead, they arise as a result of higher-order, multibody interactions which emerge naturally as the coupling strength is increased above the small coupling limit. Hence the conclusion: exotic states can arise even in simple networks which have linear, pairwise, nearest-neighbor physical coupling.

2.6 Conclusion

We introduced a variety of canonical models of coupled oscillators and highlighted the key features of each which are relevant to this work. Then, we introduced the amplitude-phase model which describes these experiments and derived its higher-order phase dynamics. We will reference these models throughout the text.

Chapter 3

COMPUTING WITH NANOMECHANICAL OSCILLATORS

In this chapter we introduce some methods for computing with coupled oscillators. We focus on two of these which we call annealing and reservoir computing. Beyond these two, associative memory has been of considerable interest and may be amenable to implementation on coupled oscillator platforms [30–43]. Many other methods for implementation and training oscillatory computers have been suggested [44–46] but are beyond the scope of this discussion; see [47] for a review. There are still many challenges and open questions on both the algorithmic side and the hardware side of the problem. To bring these methods to life, particularly for commercial applications, there is a substantial amount of research left to do.

3.1 Background

Most modern-day digital computers operate using the von Neumann architecture. In this architecture, the processor and memory are separated such that on each computational cycle, the processor must fetch the data from memory, perform the desired operation on that data, and then store the result back to memory. This occurs sequentially on one or many shared buses, creating what is known as the von Neumann bottleneck. This approach has proved to be remarkably powerful for building general-purpose computing machines. But, it is inefficient due to its inherently sequential operations. However, computers have continually improved because the underlying hardware, i.e., the number of transistors in a processor and the speed at which they operate, has exponentially improved over time. In recent decades, those trends have slowed or halted [48], leading companies and researchers to search for performance improvements by other means.

Without turning away from the CMOS-digital paradigm, a number of improvements can still be found in hardware, software, and algorithms, as reviewed in [49]. In hardware, modern central processing units (CPUs) offer a high number of cores, multiple memory channels, and advanced execution routines, alleviating the communication bottleneck. Additional specialized processors such as graphics processing units (GPUs) or tensor processing units (TPUs) have been developed which are inherently more parallel and enable performance boosts for common operations, particularly for linear algebra. These developments represent the beginning of the

transition away from general-purpose computing hardware and towards specialized processors which are faster and more energy efficient.

In addition to the specialization of CMOS-digital processors, there has also been renewed interest in entirely novel computing methods based on new types of devices and novel computing architectures. This is driven not only by the stagnation of CMOS-digital technology, but also by shifting computing demands, such as the explosion of energy consumption by large language models [50] or the demand for computing at the edge [51]. Under this research umbrella, there are many topics ranging from building logic circuits with low-power and high-speed analog devices to in-memory computing with exotic materials. These schemes fall under many categorizations in the literature, such as neuromorphic computing, brain inspired computing, thermodynamic computing, and physical computing. In the next section, we provide a few factors that must be considered to build such a computer.

Considerations for Analog Computing

For an analog computing system to be useful, it must surpass the existing technology — principally CMOS-digital processors — in one or more metrics, such as speed, power consumption, accuracy, bit precision, footprint, security, cost, versatility, or robustness.

One important consideration for special purpose analog hardware is that the need for high precision must be relaxed. Analog hardware will likely never compete with the allowable precision of modern digital processors, but there are many applications for which low precision or stochastic computations are acceptable or even superior. Successful analog computing approaches will likely focus on low precision computations or tasks which may benefit from thermodynamic fluctuations as opposed to suppressing them. For the case of optical computing, [52] provides a thorough review of such considerations and how analog optical computing may rise to a practical technology.

In terms of computational speed, computing tasks can be roughly divided into two categories: tasks where we want the answer as fast as possible, such as optimization, scientific computing, and communication, and tasks where there is a natural timescale to the problem and exceeding that timescale garners no additional benefit, such as feedback control, speech and image recognition, and prediction. To give a concrete example, say you want to turn on your TV with a voice command. There is a limit on the frequency content of the human voice and it takes a certain amount

of time to say a voice command. Thus, you do not need to sample and process the data at GHz frequencies if you have a highly efficient and parallelized voice recognition system which operates only sufficiently fast. In fact, the human brain is very slow compared to any modern computer. Timescales for sensory perception and cognition range from the exceedingly slow, around 0.5 Hz for motor function, to as fast as 20 kHz for processing directional hearing cues [53]. These slow operating speeds are compensated by the approximately $\sim 10^{11}$ neurons which are interconnected by $\sim 10^{15}$ synapses [54], all working in parallel with integrated memory and processing. The human brain may not be optimal for all computing and control tasks, but it is highly efficient for computing in its surrounding environment and exemplifies how relatively slow networks can efficiently process information via parallel in-memory computation.

3.2 Annealing for Combinatorial Optimization

Combinatorial optimization describes a class of problems where we seek an optimal solution out of an exponentially large set of possibilities. Many of these problems are in the NP-hard complexity class, and there is no known algorithm which can solve them in polynomial time. Only a subset of them, the NP-complete class, have solutions which can be verified in polynomial time. While it has not been proven as of this writing, it is generally believed that $P \neq NP$ and that no polynomial time algorithms will be found to solve NP-hard problems.

A common example is the traveling salesman problem, where we seek to find the minimum length path between a set of N locations which have an arbitrary spatial layout — for example, visiting the capital city of every country in the world. There are also many examples formulated as graph partitioning problems, such as MAX-CUT. In this case, we have a graph, $G(V, E)$, with V vertices and E edges. We wish to divide ("cut") the graph into two groups such that the maximum number of edges cross between the two groups. These problems can be binary or weighted and have many practical applications ranging from vehicle routing and logistics to resource allocation. They also show up in fundamental physics and computer science. However, as the problem size grows, von Neumann computers are not well suited for tackling their exponential complexity. Thus, we seek alternative algorithms or computing hardware which will aid in finding high-quality solutions to these types of problems.

Many of these NP-hard problems can be formulated as finding the ground state of the Ising Hamiltonian after a mapping which incurs no worse than a polynomial complexity penalty [55]. The Ising Hamiltonian with an external field is given by

$$H = -\sum_{(n,m)\in E} J_{nm}\, s_n s_m - \sum_i h_n s_n \tag{3.1}$$

where s_n is the n^{th} spin, J_{nm} is the ferromagnetic or anti-ferromagnetic coupling between the spins, and h_n is an optional external field. There are many classical algorithms which attempt to find the set of spins which minimizes the energy, H. Alternatively, researchers are working on a number of technology platforms which harness the intrinsic physics of an engineered system to find low energy states of this Hamiltonian for a given J_{nm}. These are called Ising machines and they rely on engineering interactions between artificial spins. There have been many demonstrations including many examples using optical parameter oscillators [56–58], superconducting circuits [59, 60], electronic oscillators (ring, LC, relaxation) [61–65], and non-equilibrium polariton condensates [66–68]. Each platform uses different physics to emulate and couple the spins; for a review of Ising machines and the various hardware platforms, see [69].

In addition to the physics, there must be a protocol or annealing schedule that has to be followed to ensure that the optimal solution or a semi-optimal solution is found as opposed to a potentially high energy local minimum. Because polariton condensates have highly similar dynamics to oscillators used in this work, we will review the method proposed by Kalinin et al. [67]. With slight modifications to their notation, they describe the physics of N spatially separated polariton condensates by the following equation

$$\frac{d\Psi_n}{dt} = \Psi_n\left(\tilde{\gamma}_n - \sigma|\Psi_n|^2\right) - iU|\Psi_n|^2\Psi_n + \sum_{\substack{m=1\\ m\neq n}}^{N} \Upsilon_{nm}\Psi_m + h_{kn}\Psi_i^{*(k-1)} + \xi_n(t) \tag{3.2}$$

where $\Psi_n = \sqrt{a_n}e^{i\phi_n}$ is the complex amplitude of the n^{th} condensate, a_n is the number density, ϕ_n is the phase, $\tilde{\gamma}_n$ is the effective gain, σ is the nonlinear dissipation, U is the in-site interaction strength, Υ_{nm} is the inter-site coupling, h_{kn} is the parametric pumping strength at the k^{th} harmonic, and $\tilde{\xi}_n$ is a Langevin noise. With the addition of coupling, this equation has a one-to-one correspondence with

the Stuart–Landau model described in (2.8) because they are both gain-dissipative systems with quadratic nonlinear in-site interactions and quadratic nonlinear dissipation. The parametric term has one-to-one correspondence with the parametrically pumped feedback oscillator (2.24). Importantly, Kalinin et al. propose adiabatically turning the system on by increasing the gain and coupling according to (3.3) and (3.4)

$$\frac{d\tilde{\gamma}_n^{\text{inj}}}{dt} = \epsilon \left(a_{\text{th}} - a_n\right) \tag{3.3}$$

$$\frac{d\Upsilon_{nm}}{dt} = \hat{\epsilon} \left(J_{nm} - \Upsilon_{nm}\right) \tag{3.4}$$

where a_{th} is the threshold number density, ϵ is the gain adjustment rate, and $\hat{\epsilon}$ is the coupling adjustment rate. By turning on the gain and coupling exponentially slowly, they show that the system finds fixed points which correspond to the ground states of the Hamiltonian given by (3.5):

$$\mathcal{H} = -\sum_{n=1}^{N} \sum_{\substack{m=1 \\ m \neq n}}^{N} J_{nm} \cos \phi_{nm} - \rho_{\text{th}}^{(k/2)-1} \sum_{n=1}^{N} h_{kn} \cos \left(k\phi_n\right) \tag{3.5}$$

where for $k = 2$ this corresponds to the Ising Hamiltonian, and for $k > 2$ it corresponds to the k-state planar Potts Hamiltonian. The important part here is that by turning on the coupling adiabatically, the system is able to find the global minimum. To guarantee the optimal solution, the system must be turned on infinitely slowly. In practice, that is not possible and the system does not always find the optimal solutions. This is one proposal for an annealing algorithm, and for more details about the experiment and modeling, see the paper [67]. Other protocols for an annealing schedule have been proposed and implemented [56, 59, 64]. They often depend on the underlying physics of the hardware being used and can strongly influence the time to solution and solution quality.

To date, an Ising machine utilizing direct feedback nanomechanical oscillators has not been demonstrated. But, it is possible that NEMS networks may be suitable to building energy efficient and versatile Ising machines in a compact footprint.

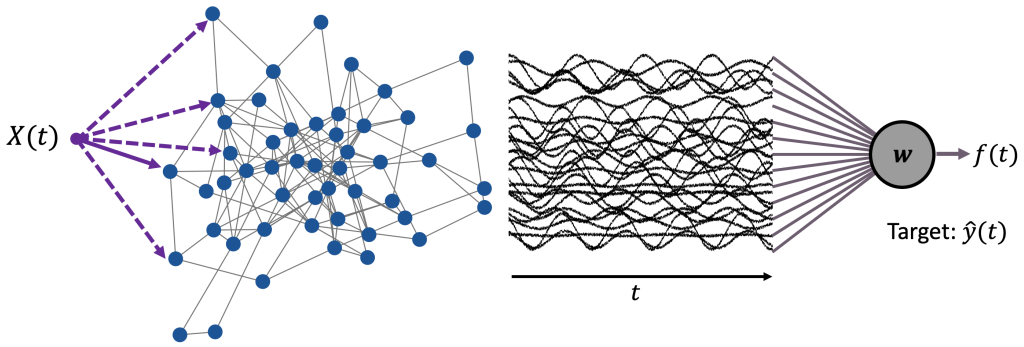

Figure 3.1: Reservoir Computing Overview. A schematic depiction of the reservoir computing concept. From left to right: the input, $X(t)$ is injected into the network via one or many nodes. The network evolves under the influence of the input. Then, a linear map, w, can be learned which minimizes the difference between the output, $f(t)$, and a target function, $\hat{y}(t)$.

3.3 Reservoir Computing

A reservoir computer (RC) is a type of recurrent neural network (RNN) where the network is composed of a fixed high-dimensional nonlinear system which maps an input into a nonlinear feature space from which a linear map can be used to learn a desired target function. The concept is illustrated in Fig. 3.1. The idea is straightforward compared to training a full RNN and amenable to physical implementations. As a result there has been a lot of interest both theoretically and experimentally in realizing practical reservoir computers to implement low-power computing and edge AI applications, among others. RC is well suited for analyzing sequential data, such as for time series prediction or speech recognition. But, it has also been implemented for image classification and other non-time-domain tasks.

Agnostic to the form of the RC, the advantages of the reservoir computing scheme are many-fold. Firstly, the reservoir is fixed, which negates the need for complex, compute intensive algorithms necessary for training the network parameters. This also reduces the hardware complexity needed to implement the scheme. In many cases, the number of required network connections scales as $O(N)$ as opposed to $O(N^2)$ in cases where all edges need to be trained. The training is also easy to implement because only the linear map at the output is trained, which can be done efficiently with regularized least squares regression. Finally, the reservoir can multitask: a single set of states from a network can be used to compute multiple output maps which compute different functions from the input data.

On the other hand, there are a few drawbacks to reservoir computing. Firstly, the hyperparameters of the network must be carefully chosen for the computer to work

well or at all. There are a number of papers describing optimal reservoir recipes [70], but there is still a lack of theoretically grounded, rigorous methods that can robustly guide practitioners in choosing performative hyperparameters. Across all aspects, researchers have yet to converge on universal principles, with one of the most prominent being whether or not RCs perform best at 'the edge of chaos' or 'edge of stability' [71–73]. Another drawback is while the reservoir is capable of multitasking, it is not optimized for any specific task. Thus, higher accuracy may be achieved with a successfully trained RNN.

For a reservoir computer to operate well, it must satisfy a few requirements [74]:

1. Nonlinearity: the network or readout must contain nonlinearity in order to faithfully reproduce strongly nonlinear mappings and classifications, as well as nonlinear dependencies in time series prediction problems.

2. Requisite complexity: The reservoir must generate a sufficiently complex set of nonlinear features from the input process. The complexity is tied to the nonlinearity and effective dimensionality of the reservoir.

3. Fading memory/echo state property: the network must forget inputs from sufficiently far in the past.

4. Separation: the reservoir must be able to differentiate inputs of different classes. But, it should still keep nearby data points together. This is closely related to the complexity.

5. Unifilarity/determinism: in general, a reservoir should evolve deterministically, but nondeterministic evolution is an ongoing topic of research.

For a more detailed review of the concepts, see [73].

We define two types of reservoir computing: conventional and physical. Conventional RC describes standard implementations of reservoir computing, typically formulated as echo state networks (ESN), which are often simulated on digital computers. Physical reservoir computing (PRC) describes the idea of using physical systems to act as the reservoir and can have widely varying technology platforms, dynamics, and architectures. Next we review these two approaches and introduce simple case studies to illustrate important concepts.

Conventional RC: The Echo State Network

The typical RC scheme [75–77] has a setup which looks like the following:

$$\mathbf{r}_{i+1} = (1-\alpha)\mathbf{r}_i + \alpha\Phi(\mathbf{A}\mathbf{r}_i + \mathbf{W}^{in}\mathbf{x}_i + \mathbf{b}) \tag{3.6}$$

$$\mathbf{y}_{i+1} = \mathbf{W}^{out}\mathbf{r}_{i+1} \tag{3.7}$$

with $\mathbf{r}_i \in \mathbb{R}^{d_x}$ being the discrete time reservoir states of dimension d_r, $\mathbf{A}$ is the $d_r \times d_r$ matrix defining the reservoir topology, $\{\mathbf{x}_i\}_{i=1}^{K} \subset \mathbb{R}^{d_x}$ is the discretized input process, $\mathbf{W}^{in}$ is a $d_r \times d_x$ linear input map, $\mathbf{b} \in \mathbb{R}^{d_r}$ is a bias vector, Φ is a nonlinear function, and α is the leak rate. Typically, $\Phi(z) = \tanh(z)$ or $\mathrm{ReLu}(z)$. $\mathbf{W}^{out}$ is the $d_x \times d_r$ linear output map.

During the training of an RC, everything ($\mathbf{A}$, $\mathbf{W}^{in}$, $\mathbf{b}$, α,...) is fixed except for the linear output map, $\mathbf{W}^{out}$. However, the fixed hyperparameters may need to be tuned for the reservoir to work optimally or even at all. How to tune these hyperparameters is a significant and open challenge in the field. Typically, the reservoir is constructed randomly. For example, in [75], the random processes defining the reservoir are

$$A_{ij} \sim U(-\beta, \beta) \tag{3.8}$$

$$W^{in}_{ij} \sim U(0, \gamma) \tag{3.9}$$

$$\mathbf{b} = \mathbf{0}. \tag{3.10}$$

The RC is used for approximating a nonlinear map, F, which takes an input, $\mathbf{X}(t)$, to an output, $\hat{\mathbf{y}}(t)$:

$$\hat{\mathbf{y}}(t) = F\big(\{\mathbf{X}(t+z)\}_{z\in[-\tau,0]}\big). \tag{3.11}$$

During the training phase, $\mathbf{W}^{out}$ is found such that

$$\mathbf{W}^{out} = \arg\min_{\mathbf{w}} \left[||\hat{\mathbf{y}}(\ell) - \mathbf{R}^{train}\mathbf{w}||_2^2 + \eta||\mathbf{w}||_2^2 \right] \tag{3.12}$$

where $\mathbf{R}^{train}$ is the discrete time reservoir states collected during the training interval, $\hat{\mathbf{y}}(\ell)$ is the discrete time sampled target, and η is the Tikhonov regularization parameter. In the test phase, we continue to evolve the reservoir according to the

dynamics and the learned mapping, $\mathbf{W}^{out}$, is used to compute the output. Such a RC has been shown to perform well on a variety of complex tasks such as chaotic timeseries prediction and handwritten digit recognition.

This method of reservoir computing is typically carried out by digital simulation, but could, in principle, be implemented in hardware. However, there are also other constructions of reservoir computing which may be more amenable to certain physical systems. In the next section, we introduce PRC for the specific case of coupled oscillators.

Physical RC with Coupled Oscillators

There are a number of examples of PRCs which utilize the dynamics of coupled oscillators. On the algorithm side, [78] used a Kuramoto-model-based reservoir to explore the effect of synchronization clusters on RC performance; [79] utilized higher-order synchronization to construct an XOR gate using a two-oscillator reservoir; and, [80] proposed a wave-based computing scheme utilizing a generic coupled phase oscillator formalism. In hardware, coupled oscillators have been used to make PRCs with spin-torque nano-oscillators [81], nonlinear micro- and nanomechanical resonators [82–84], CMOS circuits, field-programmable gate arrays (FPGAs), and photonic systems [85], among others [74]. Most physical reservoirs can be categorized as continuous media, discrete networks, or single nodes with delayed feedback. In the latter case, the reservoir is composed of a virtual network. Some examples of single-node RC systems can be found in [82–84, 86–89]. Alternatively, perhaps a real network could be constructed using multiple modes of a single device [90].

Of particular interest for this work is a method of RC introduced by Choi and Kim in 2019 [24] which utilizes the transient dynamics of an explosive Kuramoto model. The method is implementable on the coupled oscillator hardware which will be described later, so we will utilize this method as a case study in RC with Kuramoto-like dynamics.

In the paper, the authors use a dynamical system which can be described by (3.13):

$$\frac{d\phi_n}{dt} = \underbrace{\omega_n + \frac{\lambda|\omega_n|}{k_n}\sum_{m=1}^{N} C_{nm}\sin(\phi_m - \phi_n)}_{\text{Reservoir Dynamics}} + \underbrace{\frac{1}{\hat{\rho}}X(\hat{\gamma}t)\delta_{n1}\sum_{l=0}^{L}\delta(t - l\Delta t_I)}_{\text{Impulse Input}}. \quad (3.13)$$

The first half of (3.13) defines the reservoir dynamics which follow that of the explosive Kuramoto model defined by (2.13), where ϕ_n is the phase of oscillator n, ω_n are the intrinsic oscillator frequencies, λ is the global coupling parameter, C_{nm} is the adjacency matrix, and k_n is the degree of node n. The second half defines how an input signal is coupled to the reservoir. Here, $\hat{\rho}$ and $\hat{\gamma}$ are the input amplitude and time scaling of the input process, $X(t)$, δ_{n1} is a Kronecker delta which maps the input solely to the first oscillator, and $\delta(t - l\Delta t_I)$ is a Dirac delta function describing the phase impulses at discrete time index, l, spaced by the interval Δt_I. We call this the impulse input method. A more complex mask function could be defined such that the input is applied to multiple oscillators with specified or random weights, but, here, we limit to the case of one input oscillator.

This RC functions by taking the input (highlighted red) and expanding it into a high-dimensional nonlinear space, $\phi_n(t)$ (highlighted blue). The computing procedure goes as follows: first, the oscillator network is synchronized with no input. Then, the network is driven by the input until the system reaches generalized synchronization, which is defined as the condition when the transient dynamics of the coupled system have decayed such that the driven network state does not depend on the initial conditions of the network. This is called the warm-up phase. Subsequently, the states of the network are recorded for 4000 input impulses. With that data, the linear map is trained to minimize the error between the output and the target function. This is called the training phase. Finally, we continue to drive the network with the input and use the state plus the linear map to compute the reservoir's output. By comparing the output with the target value, the prediction error can be computed. This is the testing phase.

One of the key features of this RC method is that the impulses excite strong transient behavior in the system. This may be why it is able to perform computations. The fixed point itself is a simple stationary state of the dynamics, but the complexity arises from the complex trajectories that the system takes while resynchronizing between impulses. The authors study the performance of this RC scheme using a network of 500 oscillators. They explore the performance over a range of parameters and tasks. Their primary conclusion is that the explosive synchronization model outperforms a non-explosive model due to the increase in criticality.

Towards the goal of PRC with nanomechanical oscillators, we replicated the scheme and studied the performance over a variety of hardware-realistic system parameters. Fig. 3.2 shows an example of the performance of a 100-node network as the global

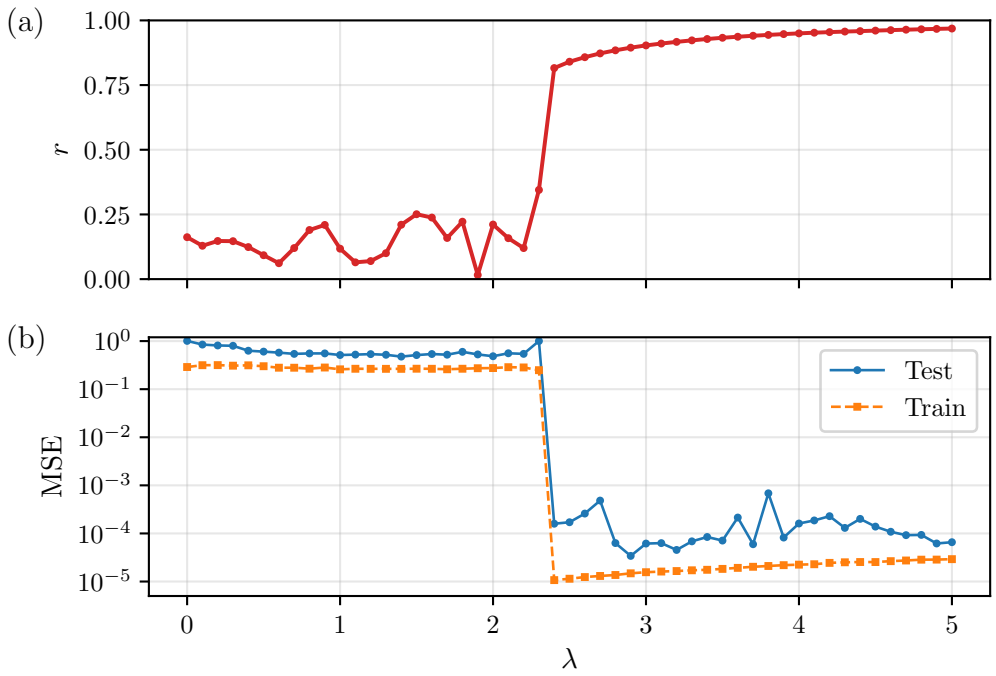

Figure 3.2: Explosive Kuramoto RC λ Sweep. Example performance of 100 node explosive Kuramoto model-based RCs with λ ranging from 0 to 5. (a) Plot of the order parameter, r, as the global coupling parameter is increased. (b) Plot of the mean square error (MSE) for training and testing of the RC for each value of λ.

coupling parameter, λ, is increased from 0 to 5. The input to the RC is the time series of the Mackey-Glass delay differential equation (DDE) and the task is to predict the value of the time series 10 units in advance. The network is constructed with the Erdős–Rényi method such that the average connectivity is $\langle k_n \rangle = 8$. The frequencies are chosen from a normal distribution with a standard deviation of 1 such that $\omega_n \sim \mathcal{N}(0, 1)$. We use 4000 training points and 1000 testing points. The linear map is trained by Tikhonov regularization with regularization parameter of 10^{-24}. Each data point represents an independent RC experiment following the described procedure.

In Fig. 3.2(a), the order parameter from the undriven reservoir, r is plotted. The sharp increase around $\lambda = 2.3$ marks the synchronization transition. In Fig. 3.2(b) the training and testing error are plotted, and a sudden improvement is observed above the synchronization threshold. We found similar results to those reported in the paper. However, in one-to-one comparison tests, we could not replicate the quantitative error values reported, and we do not know what the source of discrepancy is. Regardless, we found that this is potentially suitable and useful algorithm for the hardware that we will introduce in this work.

3.4 Conclusion

Here we introduced some concepts in computing with coupled oscillators and explored two specific approaches from the literature: Ising machines and reservoir computing. While not yet achieved in this thesis, we aim to map the dynamics of coupled nanomechanical oscillators to these physical algorithms.

Chapter 4

PIEZOELECTRIC, FLEXURAL-MODE NEMS RESONATORS

At the heart of the proposed network of nonlinear oscillators is a nonlinear resonator. The linear and nonlinear vibrations of the device determine the resulting network dynamics and must be carefully understood and controlled. With this in mind, we developed a wafer-scale fabrication method for two-port piezoelectric membrane resonators based on thin-film aluminum nitride (AlN). The devices have high quality factor (Q), strong stiffening nonlinearity, and large frequency tuning via DC bias. In this chapter, we overview the design, fabrication, and characterization of individual resonators and examine statistics from 375 measured devices. Two SEM images and the material stack of the finalized devices are shown in Fig. 4.1.

4.1 Background

To build intuition about flexural mode mechanical resonators, we start by deriving the linear and nonlinear properties of a 1D beam from the continuum mechanical model. Then, we overview the 2D case of a membrane clamped on all edges. Finally, we introduce two frameworks for understanding the full electromechanical model of the device: the modified Butterworth-Van Dyke (mBVD) model and the

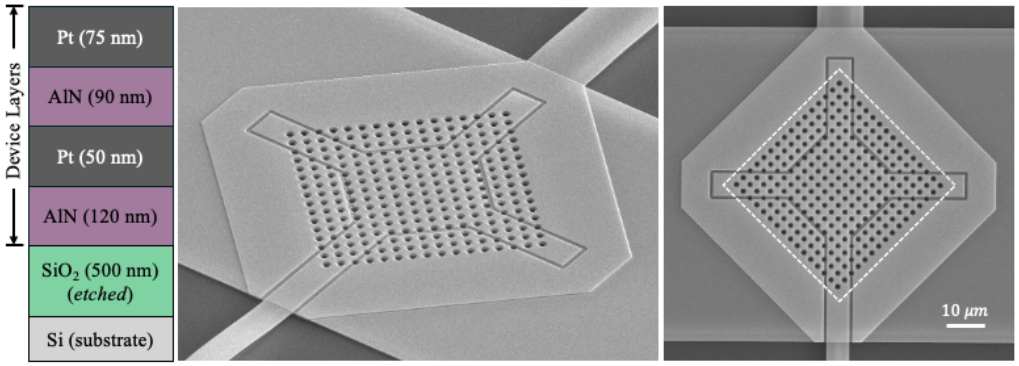

Figure 4.1: Device SEM and Material Stack. (left) material stack of the device, consisting of a seed AlN layer, a bottom Pt layer, an active AlN layer, and a top Pt layer, all on top of a sacrificial SiO_2 layer on a Si handle. (middle) Angled (52°) and (right) orthogonal SEM images of the membrane resonator. A white dashed square represents the undercut area after release.

responsivity model. For a detailed treatment of these topics, the reader is referred to [91–93].

1D Vibrations of Clamped-Clamped Beams

Following [94], the continuum mechanical equation of motion for a doubly clamped beam with Young's modulus, E, second moment of inertia, I, pre-tension, T_0, cross-sectional area, A, length, L, and density, ρ, is

$$EI\nabla^2\nabla^2 Z - \left(T_0 + \frac{EA}{2L}\int_0^L Z^2 dx\right)\nabla^2 Z + \rho A\frac{\partial^2 Z}{\partial t^2} = 0 \tag{4.1}$$

where $Z(x,t)$ is the displacement of the beam in the z direction as a function of position along the beam, x, and time, t. $\nabla = \partial/\partial x$ is the gradient with respect to x. The integral term added to T_0 is an arc-length (amplitude) dependent tension term [29]. The boundary conditions for a doubly-clamped beam are $Z(0) = Z(L) = \nabla Z(0) = \nabla Z(L) = 0$. To solve this equation, we apply the separation of variables ansatz:

$$Z(x,t) = \sum_n \Phi_n(x) z_n(t) \tag{4.2}$$

where $\Phi_n(x)$ is the shape and $z_n(t)$ is the time dependence of mode n. If we consider only the first mode, $n = 1$, and assume a mode shape of the form

$$\Phi_1(x) = 1 - \cos\left(\frac{2\pi x}{L}\right) \tag{4.3}$$

we can integrate out the spatial dependence, giving an ordinary differential equation for the time dependent amplitude of the mode [94]:

$$\ddot{z}_1(t) + \omega_0^2 z_1(t) + \alpha z_1^3(t) = 0 \tag{4.4}$$

where

$$\omega_0 = \frac{4\pi^2}{L^2}\sqrt{\frac{EI}{3\rho A}\left(1 + \frac{L^2 T_0}{4\pi^2 EI}\right)} \tag{4.5}$$

and

$$\alpha = \frac{E}{18\rho}\left(\frac{2\pi}{L}\right)^4. \tag{4.6}$$

To make the solution realistic, we add a phenomenological damping term, $(\omega_0/Q)\dot{z}$ to the equation of motion. From there, we can calculate the critical amplitude at onset of nonlinearity [95] for a beam with rectangular cross-section:

$$a_c = \frac{2}{\sqrt[4]{3}}\sqrt{\frac{1}{Q}\left(\frac{d^2}{3} + \frac{T_0 L^2}{\pi^2 E h d}\right)} \tag{4.7}$$

where d is the thickness of the beam in the direction of vibration and h is the width transverse to the vibration.

Here, we have examined the general case of a pre-tensioned beam. However, in the case that the pre-tension is very large, the restoring force due to bending can be ignored. In this case, it is called a string resonator. From (4.7), we can see that if T_0 is large, the critical amplitude at the onset of nonlinearity becomes large. Thus, in string resonators, the nonlinearity tends to be very weak. In contrast, thin beams tend to exhibit strong nonlinearities. Ref. [94] examines the onset of nonlinearity and resulting dynamic range in multiple types of beams, such as single wall carbon nanotubes and silicon carbide resonators. Typically, low onset of nonlinearity is seen as detrimental because it reduces the linear dynamic range of a device. However, low amplitude at onset of nonlinearity may be an advantage of nanomechanics for building low-power analog computers. Many of the terms in (4.7) are not independent, so it is hard to draw absolute conclusions and each case must be analyzed independently. However, this equation serves as a guide for designing nonlinear nanomechanical resonators.

2D Vibrations of Thin Plates

In this work, we utilized membranes which are clamped on all edges. The continuum mechanical equation for vibrations in the 1D case (4.1) is readily extended to the 2D case. Following [93], it can be written

$$D_p \nabla^2 \nabla^2 Z - T\nabla^2 Z + \rho h \frac{\partial^2 Z}{\partial t^2} = 0 \tag{4.8}$$

where $Z \rightarrow Z(x, y, t)$ is now the out of plane displacement as a function of x, y, and t. Here, $\nabla^2 = \partial^2/\partial x^2 + \partial^2/\partial y^2$ is the Laplacian operator. There is still a contribution

from bending, a contribution from tension, and an inertial term. However, EI in the bending term is replaced with D_P, which is the flexural rigidity of a plate, given by

$$D_P = \frac{Eh^3}{12(1-\nu^2)} \tag{4.9}$$

where ν is the Poisson's ratio of the material. The 2D clamped boundary conditions are given by

$$Z(x,y)\Big|_{x=0,L} = Z(x,y)\Big|_{y=0,L} = 0 \tag{4.10}$$

$$\frac{\partial Z}{\partial x}\Big|_{x=0,L} = \frac{\partial Z}{\partial y}\Big|_{y=0,L} = 0. \tag{4.11}$$

The devices used in this work fall in the intermediate regime where both bending and tension are relevant, but, for brevity, we refer to the devices as membranes.

The nonlinear coefficient for the first mode of a square membrane is given by [96]

$$\alpha = \frac{3E}{16\rho}\left(\frac{2\pi}{L}\right)^4 \tag{4.12}$$

which scales similarly to a beam. In practice, solving (4.8) with realistic device geometry is impractical. We use the analytical solution as a guide, but ultimately utilize finite element modeling for the design of the devices.

Beyond the analytical analysis, in this work we used 2D resonators for multiple reasons. First, when releasing a beam via isotropic etching, there is typically a ledge leftover on the support which perturbs the primary flexural modes and introduces spurious modes. This is undesirable, and, as will be seen, is circumvented with membranes. Furthermore, as derived in [92], a membrane of length L is stiffer than a beam of the same length and is expected to have a larger electromechanical coupling which means less insertion loss at the same frequency. Finally, membranes have already been successfully demonstrated in similar past experiments [27], so we follow that method.

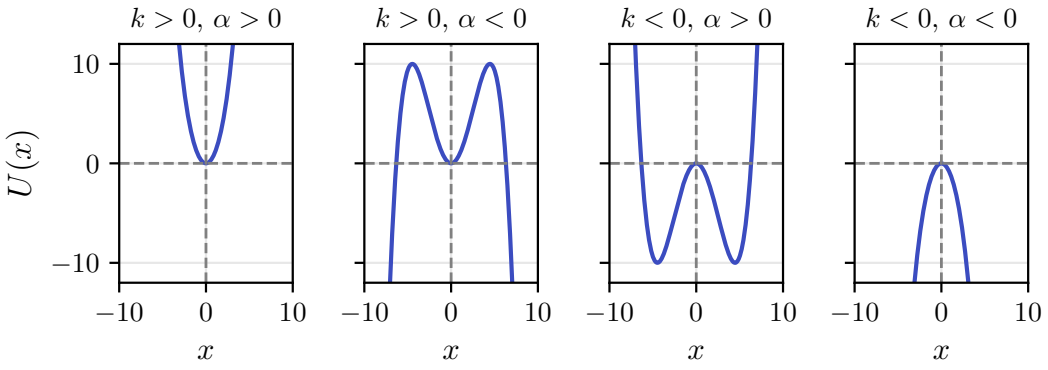

Figure 4.2: Nonlinear Potential. Example plots of the potential landscape for example parameters $|k| = 2$, $|\alpha| = 0.1$. The sign of each parameter is indicated in the title above each plot.

Variations of the Nonlinear Potential

From the lumped element model (4.4), we can see that the device displacement behaves like a particle in a potential with the form

$$U = \frac{1}{2}kx^2 + \frac{1}{4}\alpha x^4. \tag{4.13}$$

This form holds if we only keep the first term in the Taylor series expansion of the nonlinear term in (4.1) and reject the odd terms by symmetry. For the devices developed in this work, $k > 0$ and $\alpha > 0$. But, there are other devices with different nonlinearities where this is not always true. For other combinations of the parameters, k and α, this potential gives characteristically different dynamics. To be clear about the scenario under consideration, we plot the corresponding potentials in Fig. 4.2 and will quickly review each case.

Case 1 ($k > 0, \alpha > 0$)**:** the system oscillates about a single stable fixed point at $x = 0$. As the amplitude of vibration increases, so does the restoring force, leading to an increase in the resonant frequency of the system. Thus, for $\alpha > 0$, we call it a stiffening nonlinearity.

Case 2 ($k > 0, \alpha < 0$)**:** assuming that α is small (or x remains small), the system oscillates around a single stable fixed point at $x = 0$. However, as x increases, the restoring force decreases, leading to a reduction in the resonant frequency, termed a softening nonlinearity. In this case, it is also possible that at large enough x, the system escapes the basin around the center fixed point and x diverges. In reality, the system cannot escape to infinite displacement, and this case would correspond to a

catastrophic failure of the device. For example, in electrostatically actuated beams, this is known as the pull-in instability.

Case 3 ($k < 0, \alpha > 0$): the negative linear stiffness leads to an instability at $x = 0$. This is called the Euler instability and arises when the beam or membrane is under sufficient compressive stress ($T_0 < 0$) to cause buckling. In the potential, the system has two stable states which are separated by an energy barrier. Such dynamics have been studied theoretically and have been engineered in nanomechanical resonators [97, 98].

Case 4 ($k < 0, \alpha < 0$): the system has no stable fixed points. We are not aware of any study where this case is physically meaningful or realized in experiment.

The resonators used in this work follow the dynamics described by the first case, $k > 0$ and $\alpha > 0$. Thus, there is no inherent bistability in the displacement of the resonator, but we will see shortly that there is bistability in the amplitude dynamics when it is driven in the nonlinear regime. These are distinct phenomena. Henceforth, case 1 is what we mean by a nonlinear resonator, and, for clarity, we will avoid referring to it as a duffing oscillator because that terminology is often associated with case 3.

Nonlinear Vibrations

Using the Galerkin method [93, 94], the continuum mechanical model introduced in the previous sections can be solved to construct a lumped element model for the driven nonlinear dynamics of the device (at least for the 1D case). We assume (and later confirm experimentally) that the dynamics of the device are accurately captured by truncating the nonlinearity after the first symmetric term and neglecting any nonlinear damping, yielding the equation for a duffing resonator

$$\tilde{m}\frac{d^2\tilde{x}}{d\tilde{t}^2} + \tilde{\gamma}\frac{d\tilde{x}}{d\tilde{t}} + \tilde{k}\tilde{x} + \tilde{\alpha}\tilde{x}^3 = \tilde{G}\cos(\omega_d t) \tag{4.14}$$

which can be nondimensionalized with the substitutions $x = \tilde{x}\sqrt{\tilde{\alpha}/\tilde{k}}$ and $t = \omega_0\tilde{t}$, where $\omega_o = \sqrt{\tilde{k}/\tilde{m}}$ is the linear natural frequency. This yields the nondimensional equation of motion

$$\ddot{x} + \frac{1}{Q}\dot{x} + x + x^3 = G\cos(\omega t) \tag{4.15}$$

where $\omega = \omega_d/\omega_0$. The dynamics of this equation are well studied in theory and experiment. Following [29, 95], we define the small parameter $\epsilon = 1/Q$ and solve (4.15) using secular perturbation theory. For more details, see the references or Appendix A. The result is a complex differential equation which describes the slow-time evolution of the driven resonator as a function of the drive strength and frequency:

$$\frac{dA}{dT} = -\frac{1}{2}A + i\frac{3}{8}|A|^2A - i\frac{1}{2}ge^{i\Omega T} \tag{4.16}$$

where $A = a(T)e^{i\phi(T)}e^{i\Omega T}$ is a complex number describing the slow time evolution of the oscillator's amplitude, a, and phase, ϕ; $T = \epsilon t$ is the slow time scale; and Ω is the detuning of the drive frequency from the linear resonance frequency, such that $\omega = 1 + \epsilon\Omega$. By splitting (4.16) into real and imaginary parts and solving for the steady-state dynamics ($da/dT = 0, d\phi/dT = 0$) [29], we can derive the amplitude and phase response as implicit functions which depend on g and Ω. The amplitude response is given by solving

$$a^2 = \frac{g}{\left(2\Omega - \frac{3}{4}a^2\right)^2 + 1} \tag{4.17}$$

for a. Following that, the phase response can be directly computed for any particular value of a:

$$\tan(\phi) = \frac{1}{2\Omega - \frac{3}{4}a^2}. \tag{4.18}$$

To visualize the behavior of (4.17) and (4.18), we numerically solve for the solutions to (4.17) using a root-finding algorithm for different values of g and Ω and then plug them into (4.18). The resulting family of curves represents the steady-state response of the system to a varying drive and is plotted in Fig. 4.3. Solid lines represent stable branches, while dashed lines represent unstable ones. We can see in both the amplitude and phase that there is a critical drive strength where the system begins to exhibit multistability. Above this critical drive, if we sweep the frequency, the system goes through two saddle node bifurcations, one at the beginning of the bistability region and one at the end.

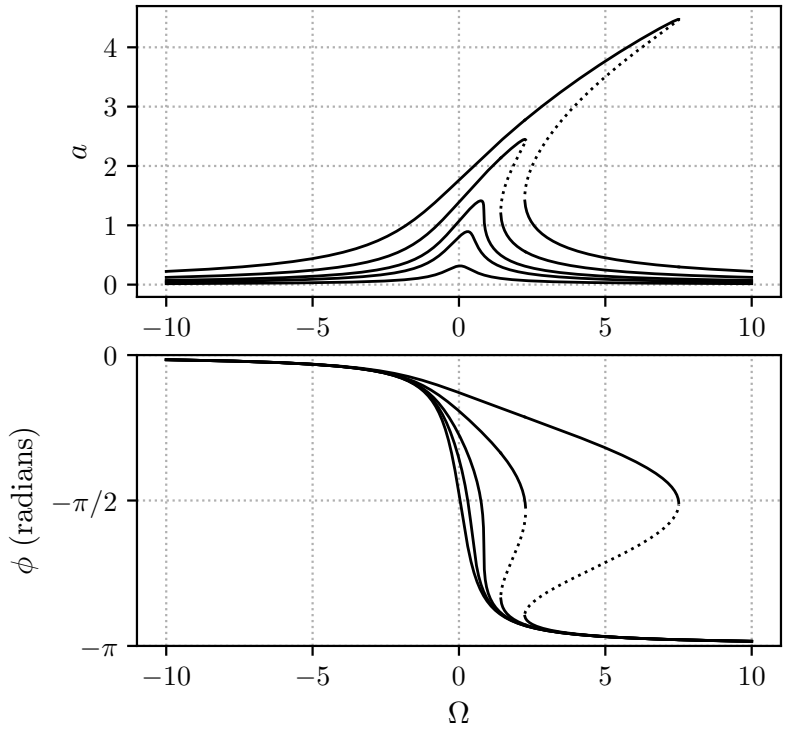

Figure 4.3: Theoretical Duffing Response. Plot of the theoretical steady-state response of a driven resonator with duffing nonlinearity, given by (4.17) and (4.18). The curves are plotted for $g = 0.1, 0.8, 2, 6, 20$. Solid lines are stable branches and dotted lines are unstable.

In experiment, this manifests as a region of hysteresis. In a quasi-static increasing frequency sweep, the device will follow the upper branch until dropping to the lower one after going through the high-frequency bifurcation point. Alternatively, in a quasi-static downward sweep, the response will follow the lower branch until it jumps to the upper branch at the lower frequency bifurcation point. The phase will follow the corresponding phase curve.

The key feature for this work is that if we connect the peak of each upper branch curve, we get a line called the backbone curve. This backbone curve reveals that the nonlinear resonance frequency, ω_r, is a quadratic function of the drive amplitude: $\omega_r \propto g^2$. Furthermore, when the drive is weak the slope of the backbone curve is nearly vertical, indicating there is very low coupling between the amplitude and resonance frequency. However, as the drive increases, the coupling between the amplitude and resonance frequency increases. This causes amplitude fluctuations to be converted into frequency fluctuations. This coupling can be useful or detrimental depending on the application.

It is important to note that this backbone curve is a fixed function of the device geometry. In a particular device, the geometric nonlinearity, $\tilde{\alpha}$, is usually not tunable.[1] However, the effective amplitude-frequency coupling can be tuned by changing the operating point of the resonator along the backbone curve.

Piezoelectric Actuation and Detection

In this section, we will overview the operating principles of flexural mode devices fabricated with thin-film AlN. For a more complete discussion of piezoelectric materials, devices, and modeling, see [92, 93, 99]. For more details about similar NEMS resonators based on AlN thin-films, including a derivation of the actuation and detection efficiency, see [100–103]. This discussion follows from those references, with adaptations for the devices used in this work.

The device is a 4-layer laminate membrane. From the bottom up, it is composed of a structural AlN layer, a bottom electrode layer, an active AlN layer, and a top electrode layer, as shown in Fig. 4.1. The top electrode is patterned to create two ports, one for actuation and one for detection. The bottom electrode covers the whole device and serves as ground. The full fabrication process is described in §4.2.

Fig. 4.4(a) shows a schematic cross-section (not to scale) of the device with each layer labeled by the material. For clarity, no release holes are depicted in this figure, but the suspended area is formed by etching the underlying SiO_2 layer, as shown. This frees the membrane to vibrate in its mechanical eigenmodes. Fig. 4.4(b) shows an expanded view of the suspended area of the device. To actuate the first eigenmode, a voltage, V_{in}, is applied to the input electrode in the center. This applies an electric field, $\vec{E}$, across the active AlN layer underneath that electrode which induces a localized polarization, $\vec{P}$, in the active piezoelectric layer. Through the reverse piezoelectric effect, this generates a strain, $\vec{\epsilon}$, both in and out of the plane via the d_{31} and d_{33} piezoelectric coefficients, respectively. The primary actuating force, however, comes from the in-plane strain resulting from the d_{31} coefficient. Due to the clamped boundary conditions, this strain results in a stress, and, importantly, that stress induces a bending moment because the active AlN layer is not symmetric about the neutral axis of the beam. Thus, by applying a time-varying voltage to the input electrode, we can effectively actuate motion in the membrane. If we apply

[1] Although, there are many demonstrations of devices with *in situ* tunable nonlinearity which have been reported in the literature.

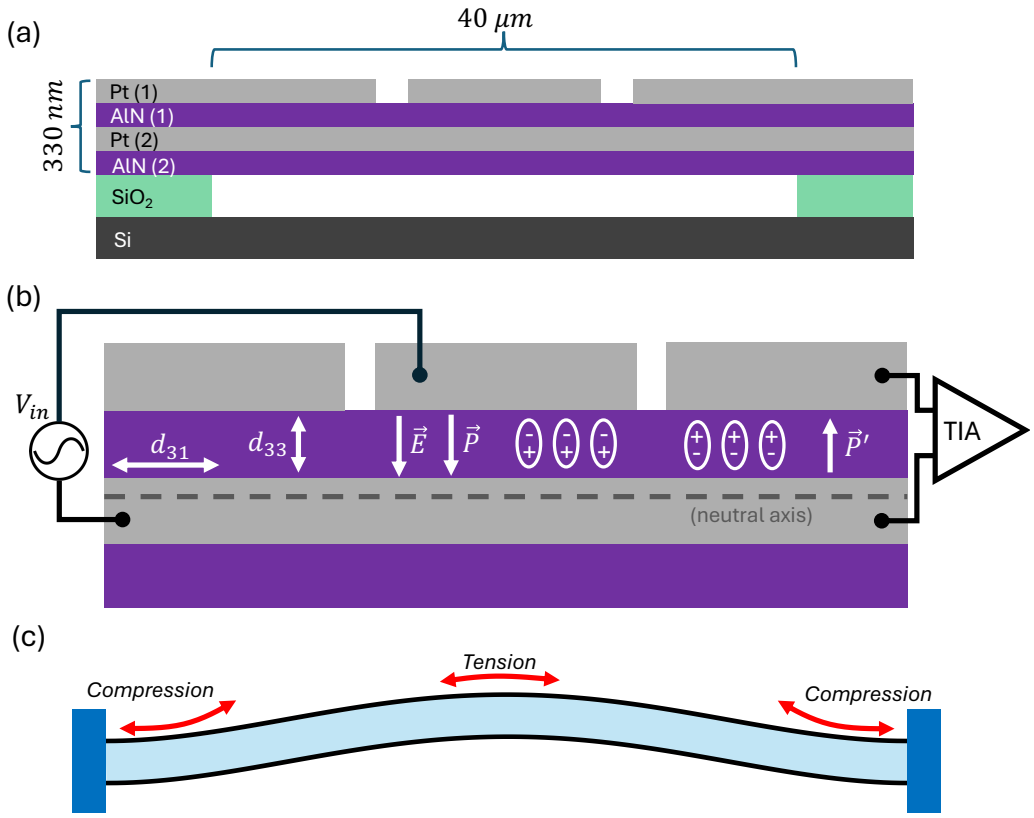

Figure 4.4: Piezoelectric Actuation and Detection. Schematic outline of the actuation and detection principles of the device. (a) a cross-sectional view (not to scale) of the device without showing the release holes. The top Pt layer shows the etched trenches which separate the input and output in the full 3D device geometry. The SiO_2 layer at the bottom is etched away to suspend the device. (b) an expanded view of the suspended part of the material stack showing the transduction physics starting from the applied voltage on the left to the transimpedance readout amplifier on the right. The neutral axis is shown to underscore the asymmetry of the active AlN layer. (c) a cartoon showing the mode shape of the first eigenmode of the membrane. Red arrows and labeling indicate the stress occurring above the neutral axis during upward deflection.

a sinusoidal voltage at a mechanical resonance frequency of the device, we can effectively actuate that mechanical mode.

The resulting motion produces a strain in the active AlN under the readout electrode. That strain induces a polarization, $\vec{P}'$, through the forward piezoelectric effect. If the device has some time-varying motion, the changing polarization creates a displacement current that can be read out with a transimpedance amplifier (TIA). The result is a two-port device which has an electrical input and electrical output, but the properties are determined by the mechanical eigenmodes. This forms an electromechanical filter which adopts the linear and nonlinear properties of the mechanics. Only electrical signals which are near the mechanical resonance frequency will be passed through the device.

In Fig. 4.4(c), a side view of the first mechanical eigenmode of a square membrane is shown. The geometry of the electrodes seen in Fig. 4.1 is designed to optimally excite and detect this first mode. As indicated, during upward deflection and for regions above the neutral axis, we expect compression near the edges and tension in the center. In downward deflection, there is compression in the center and tension on the edges. In designing the size and extent of the electrodes, the target mode for actuation and detection must be considered.

This actuation scheme enables direct, bipolar tuning of the linear resonance frequency by applying a DC bias voltage to the input electrode. Importantly, the DC voltage directly affects the stiffness of the device, allowing for strong tunability with minimal secondary effects [101]. This is in contrast to other tuning methods, such as electrostatic or electrothermal where the tuning is not bipolar and often strongly affects other properties of the resonator or the transduction.

Framework 1: The BVD Model

It is common to model electromechanical resonators using a lumped-element equivalent circuit model called the Butterworth-Van Dyke (BVD) model or modified Butterworth-Van Dyke (mBVD) model. The BVD model has two parts: (1) a mechanical resonance branch consisting of a resistance R_m, a capacitance, C_m, and an inductance, L_m; and (2) an electrical shunt branch consisting of a capacitance between input and output, C_0. The mBVD model includes two additional terms, R_0 to model dielectric loss and R_s to model resistive losses in the contacts. Fig. 4.5 shows the mBVD model with R_s being split into its input and output components,

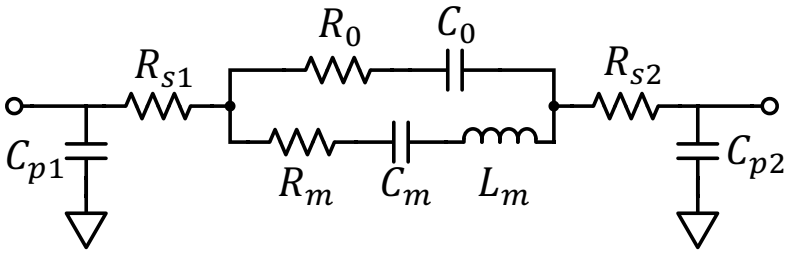

Figure 4.5: mBVD Model. mBVD electrical circuit model with parasitic capacitance. R_m, C_m, and L_m make up the motional branch. R_0 and C_0 make up the shunt branch. R_{s1} and R_{s2} are the input and output electrode resistances, respectively. C_{p1} and C_{p2} are the input and output parasitic capacitance, respectively.

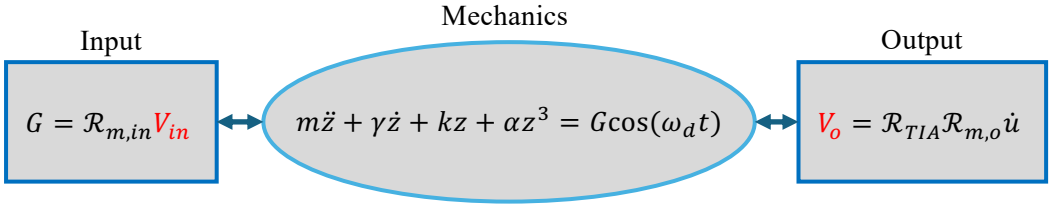

Figure 4.6: Responsivity Model. Schematic representation of the responsivity framework. The left box represents the coupling between the electrical input and resulting mechanical force. The middle is the nonlinear dynamics of the mechanical mode. The right box is the conversion from the device's velocity to output voltage, including the electromechanical responsivity, $R_{m,o}$ and the responsivity of the TIA.

R_{s1} and R_{s2}. The figure also includes parasitic capacitance terms on the input, C_{p1}, and output, C_{p2}. For a detailed analysis of this model, see [92].

The advantage of using an impedance model is that it can be directly fit to network analyzer measurements done in the lab and used in electrical circuit modeling. This is particularly advantageous for filtering and oscillator applications where the extracted electrical parameters can be used to design complex filters, interfacing electronics, or feedback loops. However, it is not straightforward to generalize the model for mechanically nonlinear devices. Thus, we introduce an alternative framework, which we call the responsivity model, in the next section.

Framework 2: The Responsivity Model

To conceptualize the electrical transmission of the devices in the nonlinear regime, we introduce the responsivity model shown in Fig. 4.6. In the schematic, the input is shown on the left, the mechanical response in the middle, and the output on the right. This model is only valid near the resonance frequency and for resonators which have

sufficiently low damping. In other words, it neglects the shunt capacitance, C_0, in the mBVD model by assuming that the motional impedance on resonance is much smaller than the impedance of the shunt branch at that frequency.

In operation, a voltage is applied to the input electrode. That voltage is transduced into force, G, by the input electromechanical responsivity, $\mathcal{R}_{m,in}$. Then, near the resonance frequency, the shunt capacitance in the electrical model (Fig. 4.5) can be ignored, leaving a standard second-order differential equation with a duffing nonlinearity. The output is a linear function of the velocity of the plate, given by the product of the output electromechanical responsivity, $\mathcal{R}_{m,o}$, and the transimpedance gain, $\mathcal{R}_{TIA}$. Altogether, we can define the transmission on resonance (forward scattering parameter) in dB as

$$S_{21}(\omega_d) = 20 \log\left(\frac{V_o}{V_{in}}\right). \tag{4.19}$$

In the linear regime, the mechanics can be solved exactly with a solution of the form

$$u(t) = U(\omega) cos(\omega t - \phi) \tag{4.20}$$

where

$$U(\omega) = \frac{G/m}{\sqrt{(\omega_n^2 - \omega^2)^2 + (\frac{\omega_n \omega}{Q})^2}}. \tag{4.21}$$

Considering only the amplitude, the velocity is given by

$$|\dot{u}(t)| = \omega U(\omega) = \frac{G\omega/m}{\sqrt{(\omega_n^2 - \omega^2)^2 + (\frac{\omega_n \omega}{Q})^2}} \tag{4.22}$$

on resonance, $\omega = \omega_r$, the response amplitude is

$$\omega_r U(\omega_r) = \frac{QG}{m\omega_r}. \tag{4.23}$$

The electromechanical responsivities are primarily a function of the mode shape, the electrode geometry, and the piezoelectric properties. The responsivity of the transimpedance amplifier is controlled by the design but can be strongly influenced by the device parasitic capacitances. In the mechanics, we generally aim to maximize

the quality factor and the nonlinear coefficient, α. These are a function of the thin-film properties, such as stress, elastic modulus, density, and quality.

Collecting the terms, the transmission on resonance can be expressed as

$$S_{21}[\mathrm{dB}] = 20\log\left(\mathcal{R}_{TIA}\mathcal{R}_{m,in}\mathcal{R}_{m,o}\frac{Q}{m\omega_r}\right). \tag{4.24}$$

Assuming that the input, output, and amplifier responsivities are linear under large drive powers, the model can be generalized to include the nonlinearity of the mechanics. In this framework, it is easier to conceptualize the contributions to the overall device performance and to consider those nonlinear effects. However, it is not necessarily easy to measure the responsivities or to use this model for peripheral circuit design.

4.2 Fabrication

The fabrication starts with a 4" $\langle 100 \rangle$ high resistivity Si wafer. On that wafer, we grow a 500 nm thermal oxide to serve as a sacrificial layer during final release. Next, we sputter 120 nm of AlN with a 0 W bias, directly followed by 50 nm of Pt sputtered at 300°C. These will henceforth be called the bottom AlN and bottom Pt layers. Using a deep ultraviolet (DUV) lithography tool and ion beam etching (IBE), we pattern the bottom Pt layer to define the bottom electrodes. Subsequently, we sputter 90 nm of blanket AlN at 6 W bias and 300 °C, immediately followed by 75 nm of Pt at the same temperature. These layers will be referred to as the active or top AlN and the top Pt. After this deposition, the top contacts are defined by DUV lithography and IBE of the top Pt. The final step is to define the release holes. We again define them with DUV lithography, followed by IBE of the entire stack down to the SiO_2 layer. The wafers are then diced down to the chip level with a dicing saw. Finally, the membrane is released by a 20 minute long wet etch in 6:1 buffered oxide etch (BOE) solution.

While fabrication is not the primary focus of this thesis, there were considerable challenges that had to be overcome to build successful devices. Briefly, some of those challenges include: (I) engineering the total thin-film stress such that the devices do not buckle. Thin-film AlN is generally a compressive material [104], particularly when deposited with a bias. That stress must be counteracted by depositing the Pt with tensile stress. (II) The layer below the AlN being deposited is very important for good film quality. Because we pattern the bottom Pt, it is a non-planar deposition.

It is not clear how this affects the film quality. (III) When etching the release holes, the top and bottom layer can easily become shorted, likely due to re-deposition of Pt on the sidewalls. This can be alleviated by etching a slightly larger hole in the top layer before etching all the way through the stack. (IV) Careful attention must be paid to the electrical characteristics of the device, including shorting between layers, parasitic capacitance, and resistance of thin-film electrodes. For this work, we made the necessary improvements to meet the project goals, but, there is still room for enhancing the performance.

4.3 Single Device Characterization

The resonant frequency, quality factor, nonlinearity, and DC tunability are the most important mechanical properties in each device. Following that, we consider the signal-to-background ratio, the electro-mechanical coupling, and the total insertion loss.

Linear and Nonlinear Properties

To characterize the devices, we perform two-port transmission measurements. The frequency is low enough that we can extract the necessary parameters for our application without an extensive calibration and without using ground-signal-ground (GSG) probes. Initial device characterization is done under vacuum ($<\sim$ 1 mTorr) using a probe station, as shown in Fig. 4.7. This method allows for rapid characterization of many devices, but it introduces higher parasitic and feedthrough capacitance than wirebonding. For the measurements, we use a Zurich Instruments HF2LI Lock-in amplifier to drive and detect the device while sweeping the frequency and power, essentially acting as a network analyzer. At the output of the device, we employ a Femto Amp HVA-200M-40-F high-impedance voltage amplifier with 40 dB gain, 1 MΩ||15 pF input impedance, and 4.5 nV/$\sqrt{\text{Hz}}$ input voltage noise. We minimize the cable length between the output and the amplifier to minimize the parasitic capacitance.

Each measurement requires three probes: input, output, and ground. For simplicity and due to availability, single-ended probes are used as shown in Fig. 4.7(c). In a typical measurement, we apply a voltage to the center electrode, detect the current on the outer electrode, and ground the bottom Pt layer. We find that the signal to background ratio is greatly improved by grounding the bottom Pt. However, the overall signal on resonance is larger without grounding the bottom Pt. This effect arises in these devices specifically because part of C_{p1} and C_{p2} (see Fig.

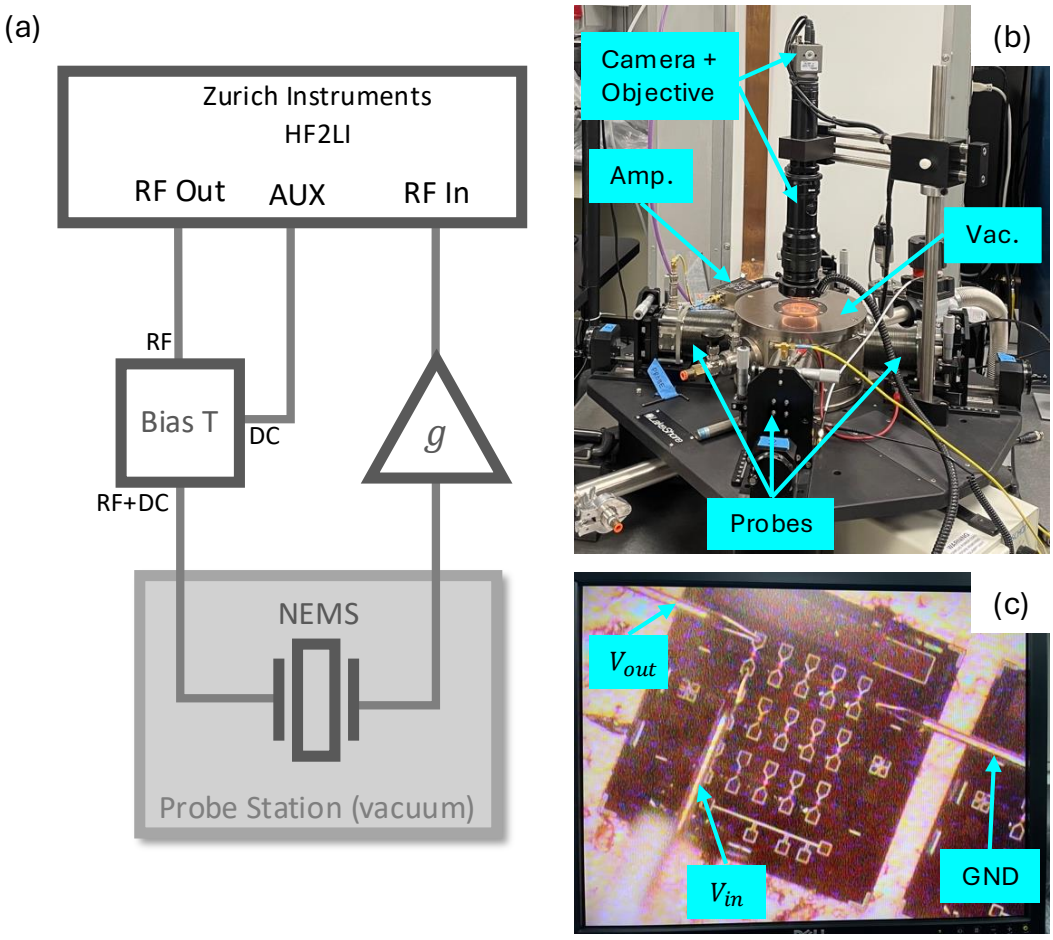

Figure 4.7: Probe Station Measurement Setup. (a) schematic depiction of the measurement showing the lock-in amplifier, DC bias, attenuation, and amplification. (b) picture of the probe station vacuum chamber. (c) image of the camera view of the device being probed.

4.5) contributes to the feedthrough capacitance when the bottom electrode is not grounded.

Here we consider only the first mode of each device. Higher modes were measured, but their characterization is outside of the scope of this document. For each device, we characterize the linear and nonlinear response and the response to DC bias. Firstly, Fig. 4.8 shows the device's amplitude and phase response for increasing drive power from -57 dBm (315 μV_{rms}) to -18.9 dBm (25.3 mV_{rms}). In this plot, the frequency is always swept from low to high. Following from the theoretical analysis in §4.1, at low drive power, we observe a typical linear resonance response. From there, we can extract the quality factor of the resonator by fitting a modified

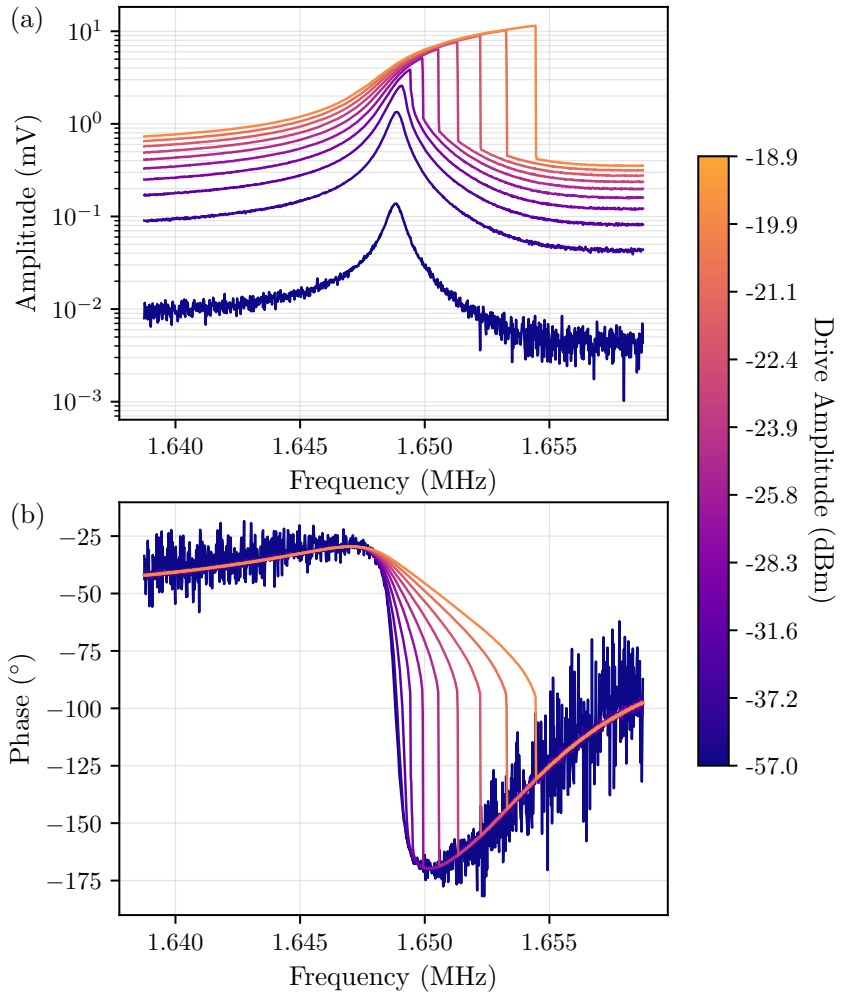

Figure 4.8: Single Device Nonlinear Frequency Response. (a) Amplitude and (b) phase response of the fundamental mode of a single device for increasing drive power. All curves are quasi-static frequency sweeps from low to high frequency. Each curve corresponds to a tick mark on the color bar, indicating the incident drive power at the input of the device.

Lorentzian to the data or by computing the linewidth at half power ($\Delta f_{3\text{dB}}$) such that $Q = f_0/\Delta f_{3\text{dB}}$, where f_0 is the linear resonance frequency.

Furthermore, in the linear regime, the phase response of the device goes from $\sim -25°$ to $\sim -175°$ as the frequency is swept through resonance, giving a total phase response of $\sim -150°$. The overall offset is due to the measurement setup, and the total phase shift through resonance is less than $-180°$ because of the shunt capacitance, C_0. The shunt capacitance also creates an anti-resonance which is not highly apparent in this device's amplitude response but is what causes the phase response to increase again at frequencies above the mechanical resonance. Additionally,

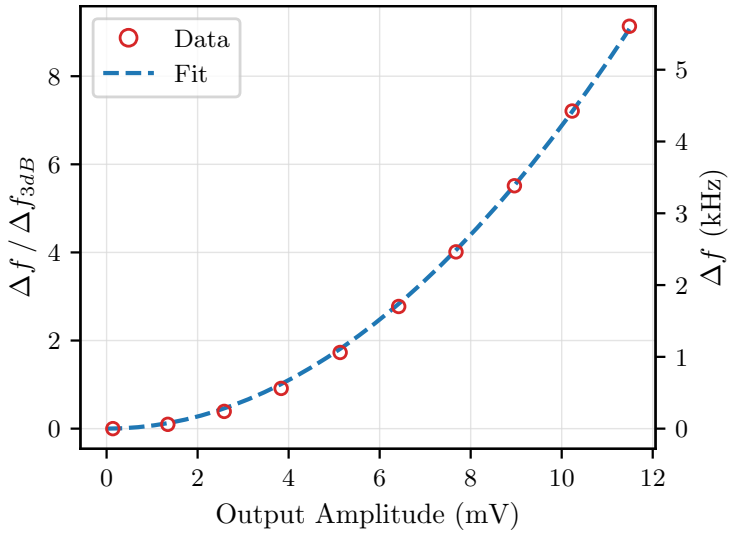

Figure 4.9: Backbone Fit. Measurement data of the nonlinear resonant frequency shift of the device as the amplitude increases. The red dots are extracted from the frequency sweeps. The blue dashed line is a quadratic fit. The left axis plots the frequency shift in linewidths, and the right axis shows the frequency shift in Hz.

the shunt capacitance sets the level of signal feedthrough at frequencies far from the mechanical resonance, called the background. By dividing the amplitude on resonance with the amplitude away from resonance, we can calculate the signal to background ratio.

As the drive power increases, the resonant frequency shifts to the right and the curve skews. Above the critical drive strength — defined as the points where the resonance curve becomes vertical — the bistability region appears. For a given drive power, the resonance response exhibits a sudden drop in amplitude at the high-frequency bifurcation point. As the drive power is further increased, the resonance frequency and the bifurcation points continue to shift upward and to the right. If the frequency was swept downward, we would observe the lower frequency bifurcation point. For each power, the phase response is also plotted and exhibits the characteristic jumps in the bistable regime.

In Fig. 4.9, we plot the extracted resonance frequency shift as a function of the drive voltage. On the left axis, it is plotted in linewidths, and on the right is plotted in kHz. We can see that even for mV drive levels, the device can be driven well into the nonlinear regime, a desirable property for low-power nonlinear devices and compatibility with standard CMOS voltage levels. From that data, we fit a quadratic function and extract the nonlinear coefficient, $\tilde{\alpha}$, in units of $1/V^2$.

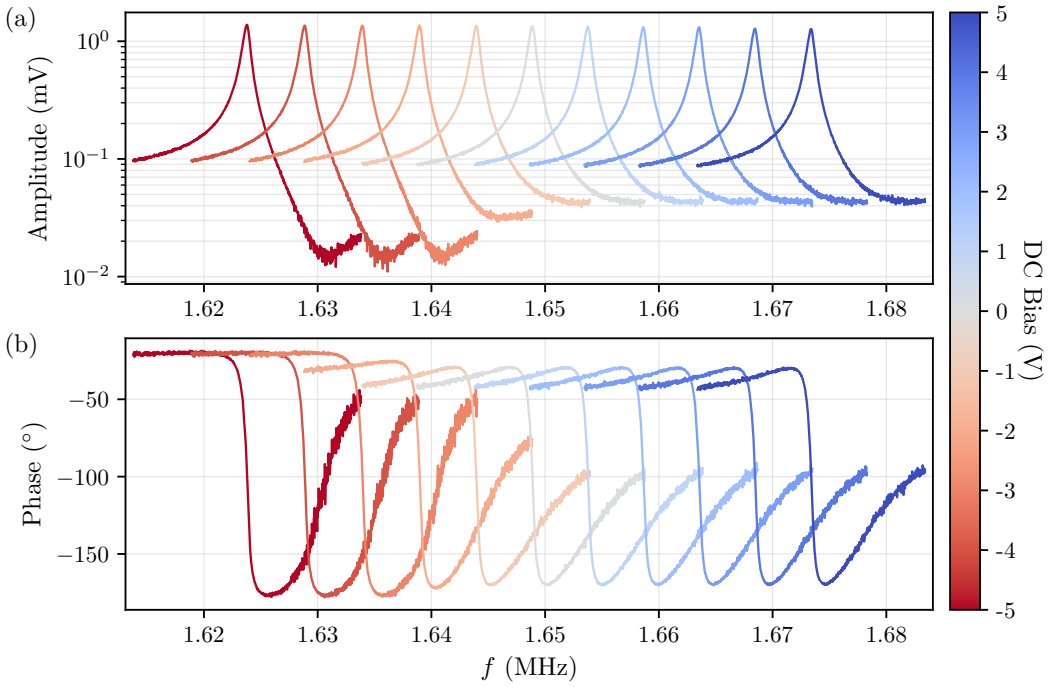

Figure 4.10: DC Tuning Frequency Response. (a) Amplitude and (b) phase response for DC bias values ranging from -5 V to +5 V in 1 V increments.

Finally, we characterize all of the resonant properties for DC bias voltages from -5 V to +5 V in 1 V increments. The full resonant responses are shown in Fig. 4.10. For each curve, we extract the linear frequency, quality factor, insertion loss on resonance, and nonlinear coefficient. These results are plotted in Fig. 4.11. Firstly, we find that the DC tuning is highly linear across the ±5 V range. We fit a line to the data and extract a tunability of 5 kHz/V (2900 ppm/V). This corresponds to a total tuning range of 50 kHz over the range, approximately 70 linewidths. We find the quality factor has a small positive trend for increasing bias, but it does not change by more than 10% across the full range. The insertion loss on resonance is also relatively unaffected, changing by less than 1 dB across the tuning range. Lastly, the nonlinearity has a clear decrease as the bias increases. This is likely due to increasing pre-tension, which increases the drive strength required to reach the onset of nonlinearity. However, more experiments are needed to rule out other factors, such as changes in the piezoelectric polarization or small pre-displacements of the membrane. Overall, the change in nonlinear coefficient is around ±9.4% from the unbiased value.

A summary of all the relevant device parameters for this prototypical device are summarized in Table 4.1. We find that all of these properties are highly consistent

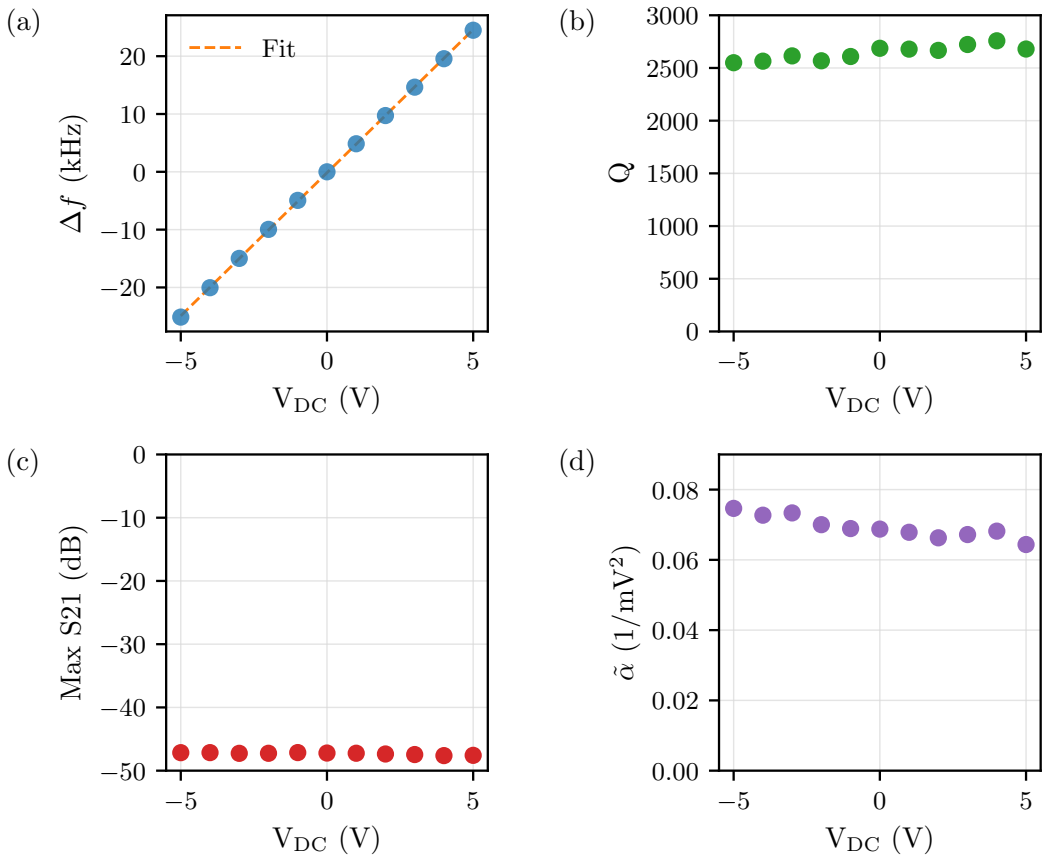

Figure 4.11: Device Properties vs DC Bias. (a) linear resonance frequency as a function of DC bias. The dashed line is a linear fit of the data (blue dots) from which the DC tunability was extracted. (b) Quality factor, Q, (c) Max S21 (insertion loss), and (d) duffing coefficient ($\tilde{\alpha}$) as a function of DC bias.

across more than 300 device measurements. Some statistics about these properties are presented in §4.4.

4.4 Ensemble Characterization

For the network studies in the upcoming chapters, we desire devices with nearly identical f_0, Q, and $\tilde{\alpha}$. However, that is not easily achievable due to intrinsic fabrication variability. Modern electron beam lithography tools allow for highly consistent geometry, even in an academic setting. But the thin-film properties, particularly of the AlN, are very difficult to control. There are stress and thickness variations across the 100 mm wafer that result in significant changes to the linear frequencies of the devices. To overcome this, we fabricate arrays of devices which have varying side length, L, and then manually select a group of devices which are

Table 4.1: Summary of Device Parameters

Parameter	Value	Unit
Linear Resonance Frequency (f_0)	1.65	MHz
Quality Factor (Q)	2560	–
Linewidth ($\Delta f_{3\text{dB}}$)	643	Hz
Time constant (τ_r)	247.5	µs
Signal-to-Background Ratio (SBR)	21	dB
Electromechanical Coupling (k^2_{eff})	0.8	%
DC Tunability	4.97	kHz/V
Nonlinear Stiffness Coefficient ($\tilde{\alpha}$)	0.069	$1/\text{mV}^2$
Drive Power at Onset of Nonlinearity	1.4	µW
Insertion Loss (IL)	47	dB

close enough in unbiased frequency that they can be tuned together by applying individual DC bias voltages.

More specifically, each chip has 15 devices with length ranging from 40 µm to 45 µm. This length variation results in a frequency spread across the chip. The resonant frequency for devices from 25 chips are plotted in Fig. 4.12(a). Each row represents a single chip, and each diamond the resonant frequency of a single device. We measured 375 devices, of which 341 had sufficiently good properties, giving an overall yield of 90.9%. However, if we remove two outlier chips which had particularly low yield (chips 3 and 5), the overall yield increases to 97.8%. The plot shows that the devices on a chip generally spread evenly, indicating local spatial consistency. However, the variation from chip to chip is larger, which suggests that there are fabrication differences between chips (likely due to release time during wet etching) and/or significant thin-film variation at cm or larger scales. More work is needed to investigate the precise cause of these trends.

In Fig. 4.12(b) the quality factor is plotted for all 341 good devices as a function of their respective frequencies. Here, each device is colored according to the chip it belongs to. We aim to demonstrate that the devices are tightly clustered in the frequency-Q space, which allows for selecting highly matched devices — a critical prerequisite for precise studies of nonlinear oscillator networks. Additional statistical plots are included in Appendix D.

These results, achieved in an academic setting, are promising indicators for future applications of flexural mode, piezoelectric resonators. Perhaps with more advanced tools and techniques, such as post-fab tuning, these devices can be made with

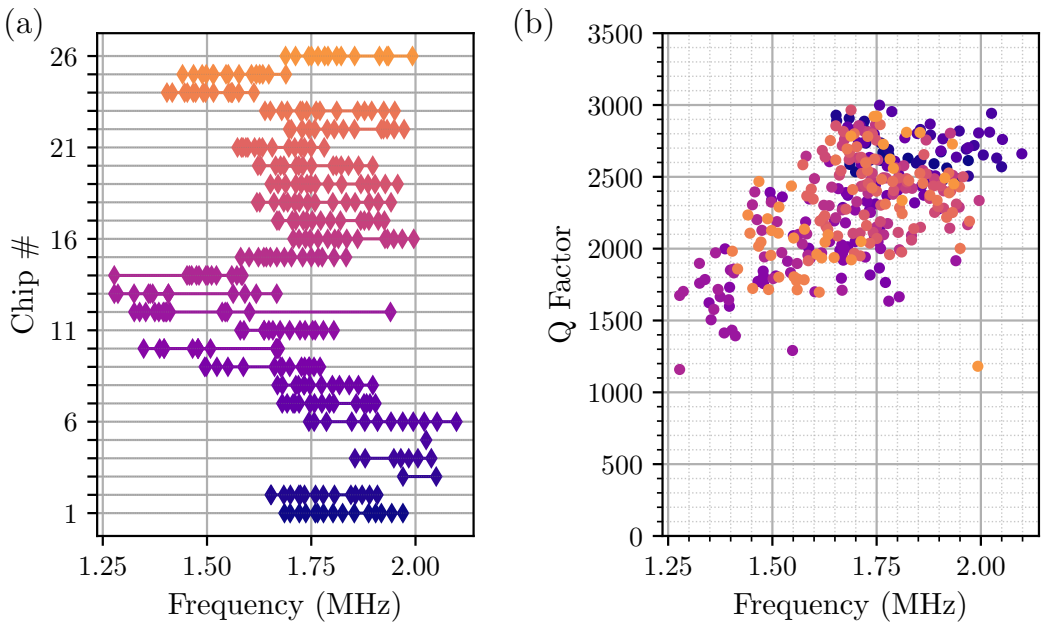

Figure 4.12: Ensemble Device Characterization. (a) Linear resonant frequency of all the measured devices. Each row represents a chip, and each diamond represents a device. It is possible that some devices are not perfectly ordered by physical dimension; that is, the frequency is not always a monotonic function of device length. (b) Q vs frequency for all 341 measured devices. The color is by chip, following plot (a).

extremely high consistency. However, this work also exemplifies the challenges associated with NEMS nanofabrication. It is likely that a substantial engineering effort would be required to achieve the repeatability and tolerances necessary for large scale commercial applications.

4.5 Conclusion

Here we have reviewed the operating principles, fabrication, and characterization of flexural mode piezoelectric NEMS membranes. We have shown that these devices are highly tunable and highly nonlinear while having very low dissipation, simple and efficient electromechanical transduction, scalability, and CMOS compatibility. These properties make them an ideal candidate technology for investigating networks of nonlinear oscillators, as we will explore in the following chapters.

Chapter 5

OSCILLATORS WITH ELECTRONIC FEEDBACK AND COUPLING

To enable experimental studies of complex networks of nonlinear oscillators, we employ electronic feedback and coupling. The electronic feedback provides gain to a resonator, allowing it to self-oscillate and transforming it from a driven-dissipative system into a gain-dissipative system. Such dynamical systems host rich behaviors that are not found in typical driven-dissipative systems, such as limit cycle oscillations and synchronization.[1] The electronic feedback is constructed from commercially available integrated circuit (IC) components which can be programmed over a wide range, enabling oscillators with highly tunable effective-nonlinearity. Similarly, electronic coupling allows for the topology, amplitude, and phase of the inter-oscillator coupling to be tuned over a wide range. Importantly, we use this method to realize complex graphs which are difficult or impossible to achieve with other coupling mechanisms.

Other studies have explored systems with mechanical coupling [107–109], electrostatic coupling [110–113], optical coupling [114, 115], and magnetic coupling [116]. These approaches are advantageous for their simplicity — they do not require additional components or power consumption to mediate the interactions. Additionally, in some cases, it is easier to engineer the desired interaction physics. For example, electrostatic force acts quasi-instantaneously compared to a typical mechanical resonator's motion, and, thus, time-delay in the coupling does not need to be accounted for. However, these coupling methods are often limited in tunability and control. In most cases, these systems are constrained to nearest-neighbor or next-nearest-neighbor coupling; or interactions where the coupling strength is proportional to the distance between nodes. Furthermore, their limited control means that fabrication variation cannot be corrected, leading to potentially undesirable parameter mismatch or asymmetry. These limitations prohibit deep exploration of dynamics on complex graphs over a wide range of parameter space.

[1] Limit cycles can be achieved in driven-dissipative systems but only by exploiting complex nonlinear phenomenon such as internal resonance or parametric excitation [105, 106].

In contrast, with electronic feedback and coupling, the amplitude and phase of every electrical path must be carefully calibrated to instantiate the correct network dynamics. In the experiment, there are multiple approaches to calibrating the setpoint of each oscillator and network edge. One such way is to break the feedback loop or coupling path at some point and directly measure the frequency response of that electrical path near the expected oscillation frequency. This works well for roughly setting the parameters, such as the phase and amplitude of the feedback. However, it fails in two ways: first, when the loop is reconnected, it can slightly change the input or output impedance of the circuits, causing unintended changes to the amplitude and phase. Secondly, as the parameters are varied for different experiments, the phase and amplitude response of the electrical paths can change. For example, the oscillator frequencies can be tuned by >10 kHz during the initial frequency corralling procedure. This large change in frequency can also affect the optimal phase setting for the feedback and coupling paths. Thus, we seek an automated calibration procedure which can be done *in situ* without breaking the loops or performing any kind of physical intervention on the hardware.

To achieve this, we turn to the known dynamics of the system. In this chapter, we introduce the architecture of the feedback and coupling electronics. Then, we derive the theoretical response of a single oscillator as the amplitude and phase of the nodal feedback and self-coupling paths are varied. Using these responses, we can identify the control parameter values commensurate with the desired network setpoints. In the final section, we review the corresponding experimental data. This chapter largely follows from previous work [27, 28, 102, 117–119] but is adapted for the system studied here.

5.1 Network Hardware Overview

Fig. 5.1 shows the circuit layout for two coupled oscillators, labeled oscillator 1 and oscillator 2. Each oscillator consists of a NEMS device embedded in a feedback loop. The key components of the feedback loop are **(I)** a multistage transimpedance amplifier, g_1, to compensate for losses in the resonator and elsewhere in the loop, **(II)** an automatic gain control (AGC) amplifier to generate saturated feedback, **(III)** a phase shifter, $\Delta\phi$, to control the phase of the feedback, **(IV)** two power splitters to couple signals into and out of the node, **(V)** a digital-to-analog converter (DAC) to apply a DC bias to the device, and **(VI)** four variable attenuators (AT1, AT2, AT3, AT4) to control the amplitude of the feedback at various points in the circuit.

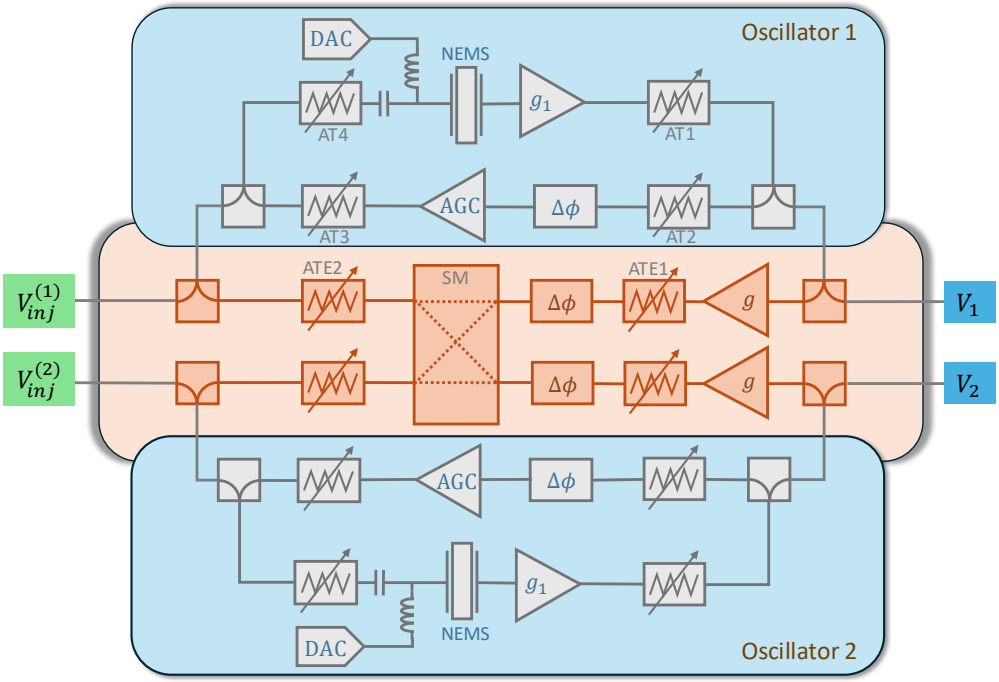

Figure 5.1: Two Oscillator Coupling Schematic. Schematic of two oscillators with switch matrix coupling. Each oscillator and its primary feedback are shown in a blue box. The coupling electronics are shown in red. Each blue box represents an individual PCB and the red represents the motherboard and network card. On the right, V_1 and V_2 are the readout signals going to the oscilloscope. On the left, $V_{inj}^{(1)}$ and $V_{inj}^{(2)}$ are external signals that can be injected into the network at each node. DAC=digital-to-analog converter, g_1=amplifier, AT=variable attenuator, $\Delta\phi$ = phase shifter, AGC = automatic gain control, and SM = switch matrix.

Of these components, only the DAC, phase shifter, and attenuators are programmable *in situ*. The other circuits can only be changed by manual replacement of passive components. The DAC controls the frequency of the oscillator by applying a DC bias to the device through a bias-T formed by an inductor and capacitor, as shown. The phase shifter is used to compensate for delays in the electrical feedback path. It allows feedback to be tuned to meet the Barkhausen phase criteria while at the resonance peak of the device, that is, the phase shift around the full loop is an integer multiple of 2π. The attenuators serve to control the operating point of the oscillator. For a detailed analysis, see [27]. In short, AT1 controls the amplitude of the output, AT2 extends the dynamic range of the AGC, AT3 controls the feedback strength, and AT4 controls the susceptibility of the node to external signals. By carefully controlling these attenuators, the properties of the oscillator can be precisely tuned.

The coupling is shown in red in Fig. 5.1. From right to left, the first element is a fixed gain amplifier (g). That is followed by a variable attenuator (ATE1) which controls the amplitude of each node's coupling output — how loud it 'talks' or its emissivity. The phase shifter ($\Delta\phi$) controls the phase of each node's output. That is followed by a switch matrix (SM) with a buffered any-to-any non-blocking architecture. This enables every edge to be independently turned off or turned on. After the switch matrix, ATE2 controls the amplitude of each node's coupling input — how hard it 'listens' or its susceptibility. Mathematically, the coupling for two oscillators is broken down as follows

$$\boldsymbol{\beta} = \begin{bmatrix} \beta_{11} & \beta_{12} \\ \beta_{21} & \beta_{22} \end{bmatrix} = \underbrace{\begin{bmatrix} \nu_1 & \nu_1 \\ \nu_2 & \nu_2 \end{bmatrix}}_{\substack{\text{Susceptibility} \\ \text{(ATE2)}}} \odot \underbrace{\begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix}}_{\substack{\text{Adjacency} \\ \text{(SM)}}} \odot \underbrace{\begin{bmatrix} e^{i\Delta\phi_1} & e^{i\Delta\phi_2} \\ e^{i\Delta\phi_1} & e^{i\Delta\phi_2} \end{bmatrix}}_{\substack{\text{Phase Shift} \\ (\Delta\phi)}} \odot \underbrace{\begin{bmatrix} \kappa_1 & \kappa_2 \\ \kappa_1 & \kappa_2 \end{bmatrix}}_{\substack{\text{Emissivity} \\ \text{(ATE1)}}} \tag{5.1}$$

where $\boldsymbol{\beta} \in \mathbb{C}_{N\times N}$ is the full coupling matrix, $\beta_{nm} \in \mathbb{C}$ is an entry in the matrix, $\kappa_n \in \mathbb{R}$ is the emissivity of node n, C denotes the adjacency matrix such that $C_{nm} \in \{0, 1\}$, and $\nu_n \in \mathbb{R}$ is the susceptibility of node n. $\odot$ signifies element-wise matrix multiplication. For example, $\beta_{12} = C_{12}\kappa_2\nu_1 e^{i\Delta\phi_2}$ is the coupling strength from node 2 to node 1. This approach allows for full control of the network's topology to realize arbitrary directed networks with control over the coupling amplitude and phase. However, control of the analog weights of the edges is constrained; ATE1 controls each column and ATE2 controls each row of the coupling matrix. The phase can only be tuned by column. Solely for the two oscillator case, the system allows for weighted, bidirectional control over the full coupling matrix.

We extend this architecture up to an eight-oscillator system. The schematic is depicted in Fig. 5.2, where one oscillator of the eight is explicitly shown. Each oscillator has an individual RF port for injecting signals, $V_{inj}^{(n)}$, and a readout port for measurement of the full nodal dynamics, V_n. The mathematical representation for the full system is given by (5.2). As in the two oscillator case, the system enables control of phase and amplitude of the columns and the amplitude of the rows, in addition to full control of the adjacency matrix.

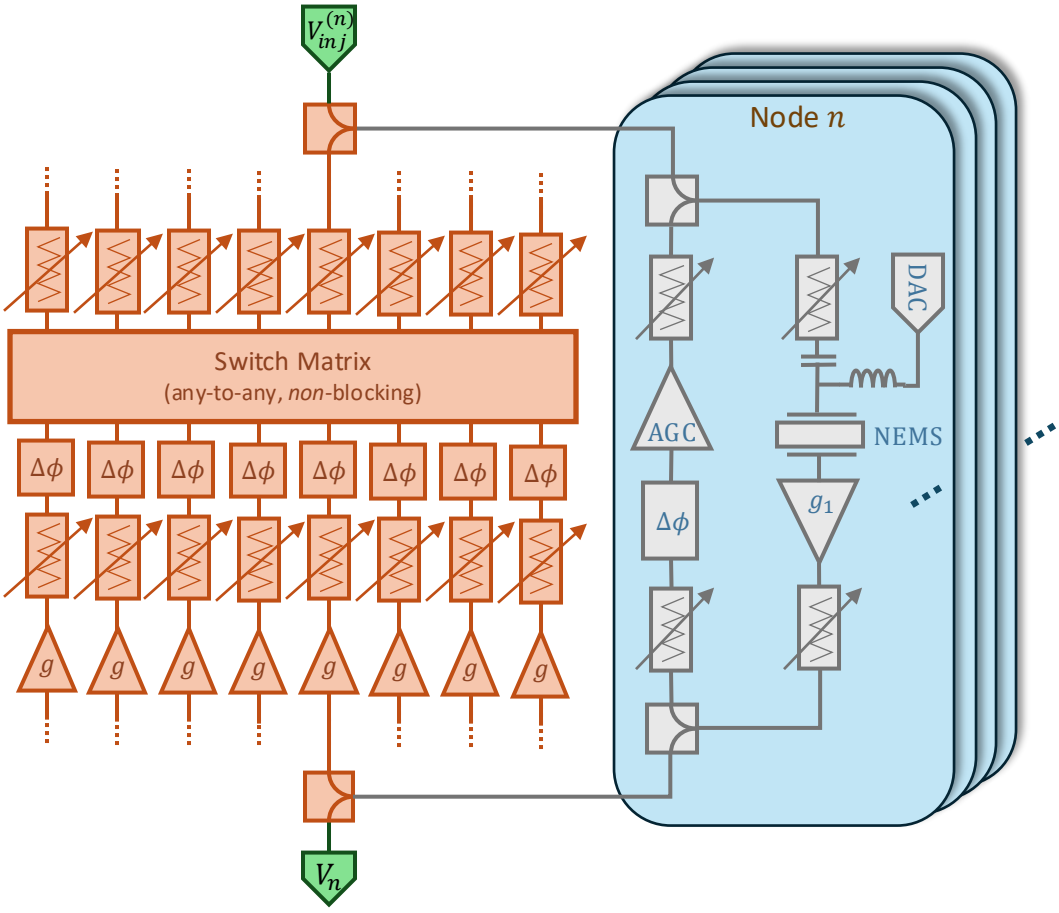

Figure 5.2: Eight Oscillator Coupling Schematic. Schematic of the network architecture. The circuit layout for the network card is shown in red, and a single oscillator is shown in blue. The network circuit is fully analog and has 8 inputs and 8 outputs. The inputs come from each node in the network, and the outputs go back to each of the nodes. The switch matrix is an 8-channel, any-to-any, buffered, non-blocking switch matrix enabling any adjacency matrix to be programmed. The system implements the coupling described by (5.2).

$$\boldsymbol{\beta} = \underbrace{\begin{bmatrix} — & \nu_1 & — \\ & \vdots & \\ — & \nu_8 & — \end{bmatrix}}_{\substack{\text{Susceptibility} \\ \text{(ATE2)}}} \odot \underbrace{\begin{bmatrix} C_{11} & \cdots & C_{18} \\ \vdots & \ddots & \vdots \\ C_{81} & \cdots & C_{88} \end{bmatrix}}_{\substack{\text{Adjacency} \\ \text{(SM)}}} \odot \underbrace{\begin{bmatrix} | & & | \\ e^{i\Delta\phi_1} & \cdots & e^{i\Delta\phi_8} \\ | & & | \end{bmatrix}}_{\substack{\text{Phase Shift} \\ (\Delta\phi)}} \odot \underbrace{\begin{bmatrix} | & & | \\ \kappa_1 & \cdots & \kappa_8 \\ | & & | \end{bmatrix}}_{\substack{\text{Emissivity} \\ \text{(ATE1)}}} \tag{5.2}$$

Variable attenuators were employed in this work, but a variable-gain amplifier could be used instead. All of the components are digitally controlled. But the oscillators and the coupling are fully analog: the signals of interest are never digitized except for readout and storage with the oscilloscope.

The slow-time equations of motion, derived in Appendix A, of the coupled oscillator network with this network card are

$$\frac{dA_n}{dT} = -\frac{i}{2}\frac{S_n A_n e^{i\xi_n}}{|A_n|} - \frac{1}{2}A_n + i\left(\omega_n + \frac{3}{8}|A_n|^2\right)A_n + \frac{1}{2}\sum_{n=1}^{N}\beta_{nm}A_m \tag{5.3}$$

where $A_n = a_n e^{i\phi_n}$ is the complex amplitude of oscillator n, S_n is the amplitude of the saturated feedback, ξ_n is the phase of the feedback, $\omega_n \in \mathbb{R}$ is the linear frequency, and $\beta_{nm} \in \boldsymbol{\beta}$ are the coupling coefficients described above.

5.2 Nonlinear Oscillators with Multiple Feedback Loops

The dynamics of a nonlinear oscillator with two feedback loops was previously derived in [102]. This subsection follows from there. For calibration, we consider only a single node in isolation. For any node, n, this uncoupled equation of motion is given by

$$\frac{dA_n}{dT} = -\frac{i}{2}\frac{S_n A_n e^{i\xi_n}}{|A_n|} - \frac{1}{2}A_n + i\left(\omega_n + \frac{3}{8}|A_n|^2\right)A_n + \frac{1}{2}\beta_{nn}A_n \tag{5.4}$$

where $\beta_{nn} \in \mathbb{C}$ is the complex self-coupling coefficient for node n. Here, the feedback saturation level, S_n, is kept as a parameter, but we will eventually scale it into the amplitude and see that it controls the effective level of nonlinearity. For the rest of the section, we consider a single node, so we will drop the subscripts n and nn on all terms for clarity. But, we emphasize that this procedure must be repeated for

every node and that the parameters can be different — intentionally or unintentionally — for different nodes. Subscripts will be reintroduced when necessary for discussing multiple nodes with parameter variation. Without subscripts, the equation becomes

$$\frac{dA}{dT} = -\frac{i}{2}\frac{SAe^{i\xi}}{|A|} - \frac{1}{2}A + i\left(\omega + \frac{3}{8}|A|^2\right)A + \frac{1}{2}\beta A. \tag{5.5}$$

From this equation, we need to calibrate the parameters of the self-feedback, S and ξ, as well as the phase and amplitude of the coupling parameter, β. We expand the feedback using $e^{i\xi} = \cos(\xi) + i\sin(\xi)$, and we let phase of the coupling vary continuously such that $\beta = |\beta|e^{i\Delta} = |\beta|\cos(\Delta) + i|\beta|\sin(\Delta)$, where $|\beta|$ is the amplitude of coupling and Δ is the phase. Plugging these into (5.5) gives

$$\begin{aligned}\frac{dA}{dT} = &\left\{\frac{S}{2|A|}\sin(\xi) - \frac{1}{2} + \frac{1}{2}|\beta|\cos(\Delta)\right\}A \\ &+ i\left\{\omega + \frac{3}{8}|A|^2 - \frac{S}{2|A|}\cos(\xi) + \frac{1}{2}|\beta|\sin(\Delta)\right\}A.\end{aligned} \tag{5.6}$$

We separate this into real and imaginary parts, which yields

$$a' = \frac{1}{2}S\sin(\xi) - \frac{1}{2}a + \frac{1}{2}a|\beta|\cos(\Delta) \tag{5.7}$$

$$\phi' = \omega + \frac{3}{8}a^2 - \frac{1}{2}\frac{S}{a}\cos(\xi) + \frac{1}{2}|\beta|\sin(\Delta). \tag{5.8}$$

From these two equations, we can identify methods to calibrate S, $|\beta|$, Δ, and ξ. Firstly, if the coupling loop is off, $|\beta| = 0$, the amplitude equation yields

$$a' = \frac{S}{2}\sin(\xi) - \frac{1}{2}a. \tag{5.9}$$

Solving for the fixed point, $a' = 0$ and letting $\Omega = \phi' - \omega$ yields the **calibration equations for the oscillator's self-feedback loop:**

$$a = S\sin(\xi) \tag{5.10}$$

$$\Omega = \frac{3}{8}S^2\sin^2(\xi) - \frac{1}{2}\cot(\xi). \tag{5.11}$$

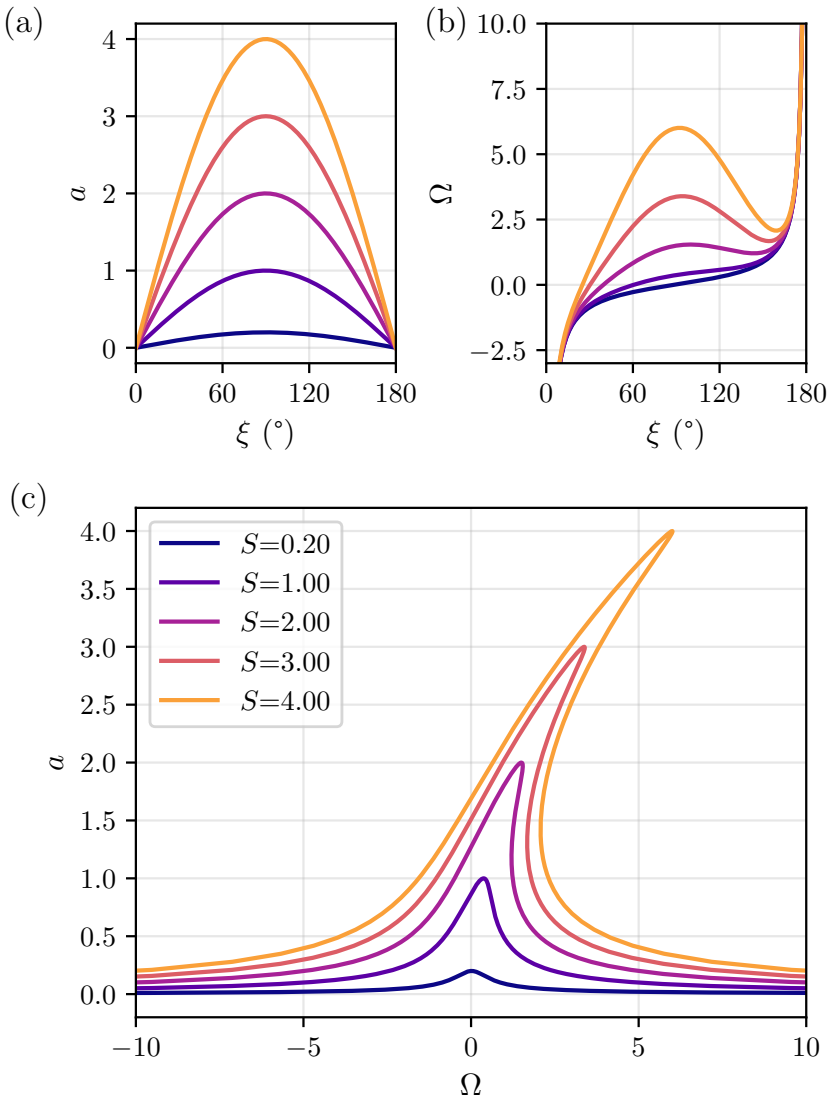

Figure 5.3: Response to Feedback Phase (Theory). Response to varying feedback phase and saturation. (a) oscillation amplitude and (b) frequency shift, Ω, as a function of feedback phase, ξ for different saturation levels. (c) Parametric plot of amplitude versus frequency from (a) and (b) as the phase is swept from 0° to 180°. All curves are stable solutions. The saturation levels for all plots are indicated in the legend in (c).

Ω is the normalized steady-state oscillation frequency relative to the linear intrinsic frequency, ω. The amplitude is maximized for any S when the phase of the feedback is $\xi = \pi/2$. Therefore, in experiment, we can find that point by sweeping the phase of the feedback and choosing the point corresponding to maximum amplitude. Fig. 5.3 shows a plot of (5.10) and (5.11) as a function of phase and a parametric plot of amplitude versus frequency as a function of the feedback phase and saturation.

From Fig. 5.3(a), we can identify that the feedback is purely dissipative when the amplitude is maximized for any value of S. It is worth noting that the frequency, Ω, is not maximized at $\Delta = \pi/2$ but slightly above that point. That point can be found be solving $\partial\Omega/\partial\xi = 0$. Here, we only plot feedback phase values up to 180°. For 180° to 360°, the amplitude is symmetric but negative, and the frequency plot is identical. In practice, the branch where the amplitude is negative means that no oscillations occur; thus, we do not consider the feedback in the range from 180° to 360°.

For the rest of the analysis, we assume purely dissipative feedback by setting $\xi = \pi/2$. Applying that to (5.7) and (5.8) yields the amplitude and phase equations

$$a' = \frac{1}{2}S - \frac{1}{2}a + \frac{1}{2}a|\beta|\cos(\Delta) \tag{5.12}$$

$$\Omega = \frac{3}{8}a^2 + \frac{1}{2}|\beta|\sin(\Delta). \tag{5.13}$$

Moreover, we scale the amplitude by S such that $a = S\bar{a}$ to further reduce the equation to

$$\bar{a}' = \frac{1}{2} - \frac{1}{2}\bar{a} + \frac{1}{2}\bar{a}|\beta|\sin(\Delta). \tag{5.14}$$

In experiment, this means that the steady-state uncoupled oscillation amplitude corresponds to $\bar{a} = 1$, regardless of the saturation setpoint. Next, we apply this scaling to the frequency equation and define the parameter $\alpha = \frac{3}{8}S^2$ to recover the form

$$\Omega = \alpha\bar{a}^2 + \frac{1}{2}|\beta|\cos(\Delta). \tag{5.15}$$

If the coupling loop is off $|\beta| = 0$, then $\Omega = \alpha\bar{a}^2$. Therefore, α is defined as the normalized frequency shift at $\bar{a} = 1$. In experiment, α is set by increasing the

feedback strength — increasing S — until the nonlinear frequency shift equals the desired value of α. The amplitude at that point becomes the amplitude corresponding to $\bar{a} = 1$.

After the self-feedback saturation is calibrated to set the desired value of α, the coupling loop is calibrated. Setting $\bar{a}' = 0$ in (5.14) and (5.15) yields the **calibration equations for the coupling loop:**

$$\bar{a} = \frac{1}{1 - |\beta| \cos(\Delta)} \tag{5.16}$$

$$\Omega = \alpha \left[\frac{1}{1 - |\beta| \cos(\Delta)} \right]^2 + \frac{1}{2} |\beta| \sin(\Delta). \tag{5.17}$$

Fig. 5.4 shows plots of (5.16) and (5.17) as a function of coupling phase and various magnitudes of β for a linear oscillator ($\alpha = 0$), including a parametric plot of amplitude versus frequency. Fig. 5.5 shows the equivalent plots for a nonlinear oscillator with $\alpha = 0.5$.

In both of the plots, four markers indicate the points corresponding to purely reactive (REAC) and purely dissipative (DISS) coupling for their attractive (ATT) and repulsive (REP) cases. These points are listed in Table 5.1 and Fig. 5.6 shows the different types of coupling by plotting β in the complex plane. Under attractive coupling, oscillators tend to synchronized in-phase, while under repulsive coupling, oscillators tend to synchronize 180° out-of-phase. This is analogous to excitatory and inhibitory interactions in neuroscience. Reactive coupling describes interactions between the displacement of two oscillators, and dissipative coupling describes interactions between the velocity of two oscillators. More details about the dynamics for different cases are described in Appendix B.

By varying the phase in the coupling loop, we can calibrate the system to these different key points in both the linear regime and the nonlinear regime. For both regimes, the dissipative points can be found by minimizing (REP) and maximizing (ATT) the amplitude.

For reactive coupling in the linear regime, the phase values can be found by finding the maximum (REP) or minimum (ATT) frequency shift. In the nonlinear case, the amplitude must be used to find the reactive points. Experimentally, we vary the phase and magnitude of the coupling concurrently and search for the phase values

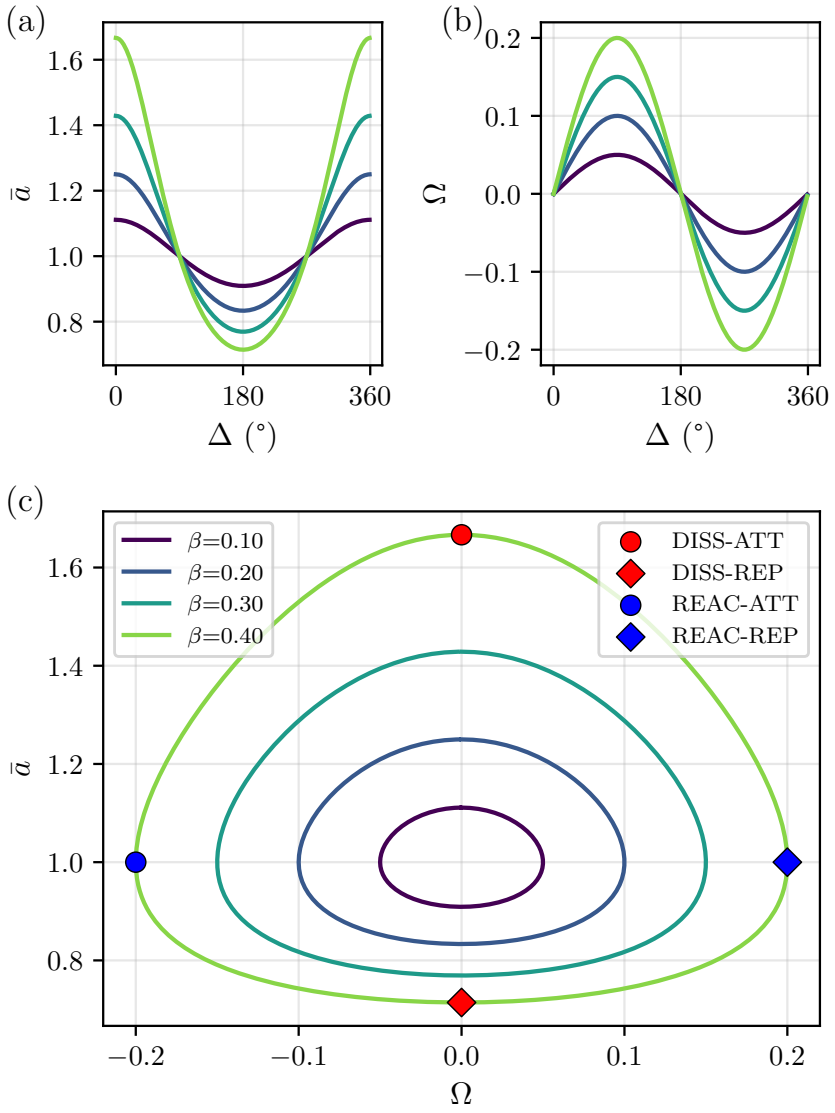

Figure 5.4: Linear Response to Coupling Phase (Theory). Response of (a) amplitude and (b) frequency of the oscillator to varying coupling loop phase (Δ) and magnitude ($|\beta|$) for a **purely linear oscillator**, $\alpha = 0$. (c) Parametric plot of the amplitude and frequency from (a) and (b), respectively, as the phase is varied from 0° to 360°. Values for $|\beta|$ for each curve in all plots are indicated in the legend in (c). Four markers in (c) indicate key values of Δ.

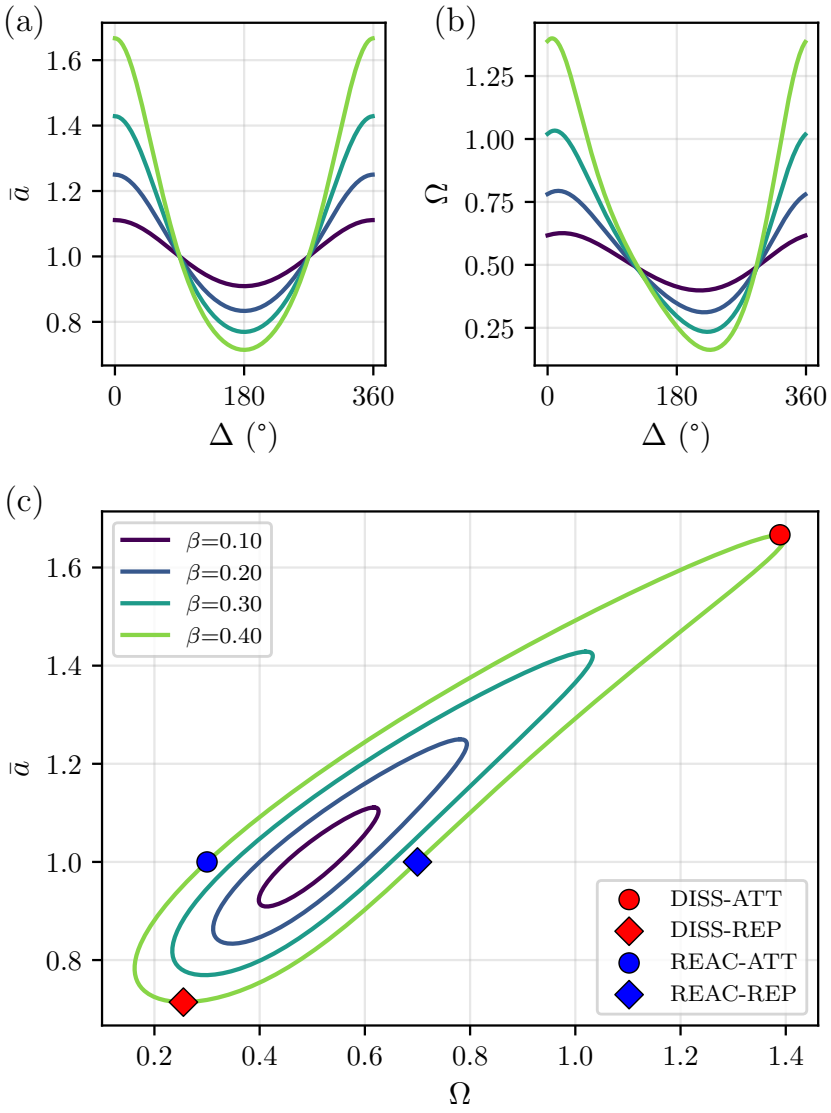

Figure 5.5: Nonlinear Response to Coupling Phase (Theory). Response of (a) amplitude and (b) frequency of the oscillator to varying coupling loop phase (Δ) and magnitude ($|\beta|$) for a **nonlinear oscillator**, $\alpha = 0.5$ in this case. (c) Parametric plot of the amplitude and frequency from (a) and (b), respectively, as the phase is varied from 0° to 360°. Values of $|\beta|$ for each curve in all plots are indicated in the legend in (c). Four markers in (c) mark key values of Δ.

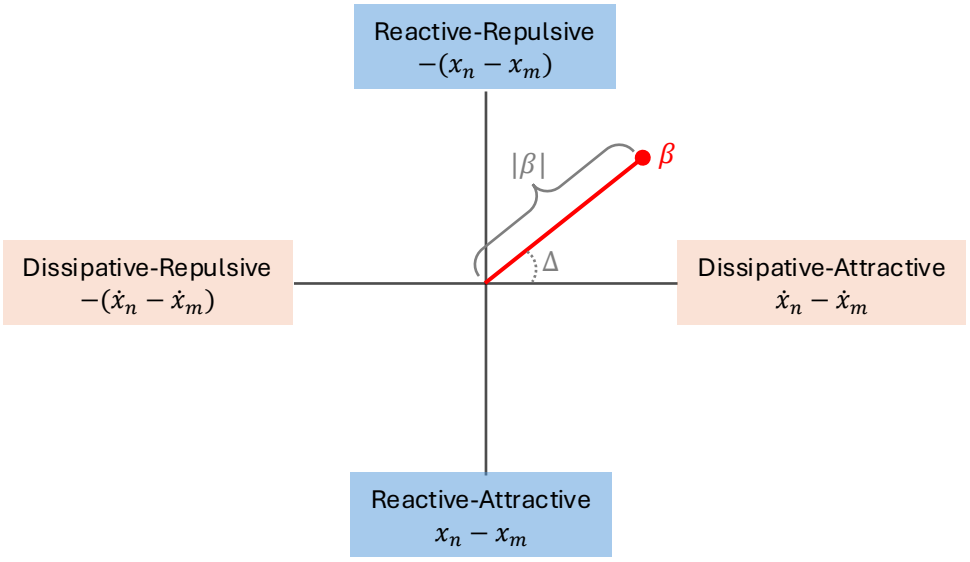

Figure 5.6: Types of Linear Coupling. β in the complex plane. Each axis is labeled by a description of the coupling type corresponding to that phase for the convention used in (5.3).

where the oscillation amplitude does not change as a function of the coupling loop strength.

Additionally, we can identify a lot of experimentally useful information from these curves. For example, it should be easier to calibrate the phase of the coupling when $|\beta|$ is larger because it will induce larger frequency and amplitude shifts, which are easier to detect in the presence of noise.

Once the phase of the coupling is set, we also wish to calibrate the magnitude, $|\beta|$. Table 5.1 shows the coupling loop calibration equations for the critical values of Δ, found by plugging Δ into 5.16 and 5.17. Experimentally, we measure the normalized frequency and amplitude shift and then calculate $|\beta|$ by taking the inverse of the equations given in Table 5.1. By using both amplitude and frequency shift data, we can confirm that the phase settings are correct.

By combining these theoretical results with experimental measurements, we can precisely calibrate the phase and amplitude of the self-feedback and the coupling paths in experiment to realize the desired interaction function electronically. The results in practice are detailed in the following sections.

Table 5.1: Equations for $\bar{a}$ and Ω at Critical Values of Δ

Δ	Type	$\bar{a}$	Ω
0°	Dissipative-Attractive	$\bar{a} = \frac{1}{1-\lvert\beta\rvert}$	$\Omega = \alpha\left[\frac{1}{1-\lvert\beta\rvert}\right]^2$
90°	Reactive-Repulsive	$\bar{a} = 1$	$\Omega = \alpha + \frac{1}{2}\lvert\beta\rvert$
180°	Dissipative-Repulsive	$\bar{a} = \frac{1}{1+\lvert\beta\rvert}$	$\Omega = \alpha\left[\frac{1}{1+\lvert\beta\rvert}\right]^2$
270°	Reactive-Attractive	$\bar{a} = 1$	$\Omega = \alpha - \frac{1}{2}\lvert\beta\rvert$

5.3 Experimental Control Parameters

In both the feedback and the coupling, we employ the same techniques for controlling the amplitude and phase of the electrical paths.

For amplitude, we utilize 7-bit digital step attenuators which are controllable from 0 to 31.75 dB of attenuation in steps of 0.25 dB. During hardware development, the attenuator performance was characterized and matched the described performance well over the full range. Importantly, the attenuation state does not significantly effect the phase shift across the electrical paths. *In situ*, open-loop characterization of the oscillator and coupling circuits found that the phase shifts by $< 0.8°$ over the full range of a single attenuator.

In the oscillator feedback loop, the AGC provides the feedback signal which has fixed amplitude at a value determined by external passive components. The AGC has a variable gain of -12 dB to +22 dB, giving it a 34 dB input range in saturation mode. Generally, this sets the range of accessible amplitude of the oscillator, but that range can be offset by programming AT2, increasing the accessible range of experimental parameters if desired. AT3 is at the output of the AGC, enabling control over the amplitude of the saturated feedback. The AGC must always operate in saturation mode, so the setpoint of an isolated oscillator should be such that the input to the AGC is far from the endpoints of the range. This ensures that amplitude fluctuations due to coupling do not drive the circuit out of saturation. Additionally, the AGC settling time must be considered if amplitude fluctuations are faster than

the AGC response, which is not the case: the nominal step response time of the AGC is 30 µs while the resonator's amplitude relaxation time is > 400 µs.

For shifting phase, we use a tunable all-pass filter that is controlled by a digital potentiometer (digiPOT). Each digiPOT can be programmed to 100 unique resistance states. In practice, the circuit allows for phase shifts ranging from ~ 0° to ~ 155° at 2 MHz. However, the phase change is not a linear function of the digiPOT setting, meaning discrete phase shift values are not equally spaced. Thus, in plots reporting sweeps of the phase, the density of points reflects both the dynamics of the system and the nonlinearity of the phase shift control curve. We do not map the phase control parameter back to actual phase shift values because it is not necessary; in all cases, we use the dynamics to calibrate the setpoints, which negates the need to exactly know the map between the phase control parameter and the actual phase shift.

One significant complication is that the phase shifting circuit does not have constant amplitude as a function of the control input. In the oscillator, the phase shifter is before the AGC, so the amplitude change can be ignored so long as the AGC is in saturation mode. For the coupling loop, changing the phase also changes the amplitude. This distorts parametric phase response curves from the theoretically expected values because we cannot easily sweep the phase while holding the amplitude constant. Despite this, we can still find the reactive coupling points by changing the amplitude and phase together, as will be described below.

5.4 Node Calibration (Experiment)

For each oscillator, we use the feedback theory derived in the previous sections to calibrate the network to the desired setpoint. Here, we review the experimental measurements used to perform this calibration in practice.

The experiments are carried out at a specific target frequency, determined by the initial spread of the oscillator's frequencies at 0 V of DC bias. Thus, we carry out all calibration procedures at or close to the desired frequency of the experiment. To do this, we tune the node we want to calibrate to the target frequency while detuning the other nodes by >20 linewidths. This is then repeated for all nodes, and enables a very precise calibration at the frequency where the final experiment will be carried out.

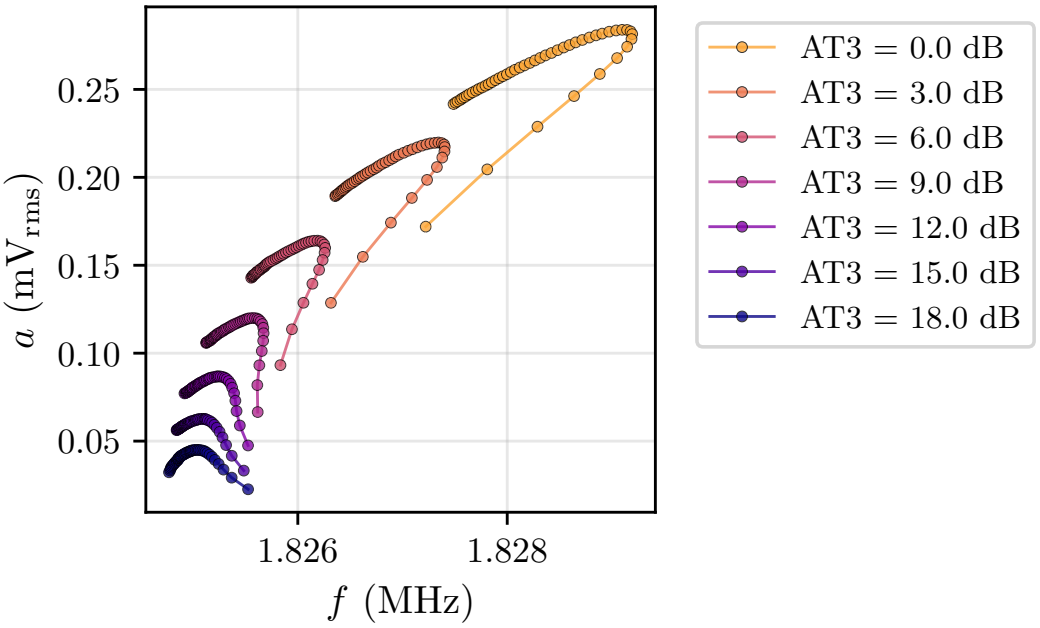

Figure 5.7: Feedback Phase Calibration (Experiment). Parametric curves for amplitude, a, and frequency of oscillation, f, as the phase and amplitude of the feedback is varied. Each curve corresponds to a particular feedback amplitude which is set by varying AT3. The phase control parameter is incremented by steps of 2 in this plot, so the density of accessible points of feedback phase is double what is shown. This data corresponds to the theoretical plots in Fig. 5.3.

Feedback Phase

In practice, the phase shift around the oscillator loop cannot be easily measured *in situ*; thus, we find the point of purely dissipative feedback by varying the phase of the feedback path and measuring the oscillator's frequency and amplitude response. In Fig. 5.7 we show the experimental data for sweeping the feedback phase at a few values feedback amplitude. As AT3 decreases, the feedback amplitude increases. We see that for low feedback strength, the oscillator has low amplitude and the phase response curve roughly tracks the shape of a linear resonance peak as the feedback phase is shifted. As the feedback strength increases, the feedback phase sweep traces out the nonlinear frequency response of the resonator. The frequency at maximum oscillation amplitude increases according to the backbone curve of the device. Importantly, we are tracking the closed-loop amplitude and phase, and, as shown theoretically, all of the branches of the oscillator's response to sweeping the feedback phase are stable. Unlike a driven nonlinear resonator, a closed-loop oscillator does not exhibit bistability, even above the critical amplitude.

From (5.10), we know that the amplitude is maximized when the feedback phase is purely dissipative. Thus, the correct phase setting can be easily determined by finding the control value corresponding to the maximum amplitude in Fig. 5.7. In the plot, each dot corresponds to an experimental setpoint. However, in the data shown, we sweep the control parameter by steps of 2 for clarity. Thus, the actual density of accessible points is double what is shown, meaning that we can precisely set the phase of the loop to realize purely dissipative feedback.

We repeat this procedure for each oscillator and recalibrate the values as necessary. In general, the values are stable over the course of days, but slow frequency and parameter drift do to temperature or other laboratory effects can cause the phase setpoints to change slightly.

Effective Nonlinearity

After correctly setting the phase of the feedback, we calibrate the amplitude of the feedback to achieve the desired effective nonlinearity, α. Experimentally, we vary AT3 while tracking the amplitude and frequency of oscillation. Fig. 5.8(a) the amplitude, a, and frequency shift, Δf, are plotted as the feedback amplitude is increased. As that happens, the oscillator climbs the backbone curve and the oscillation frequency increases. Each dot represents a setpoint of the control parameter for AT3. The frequency shift is calculated as $\Delta f = f - f_0$, where f is the actual resonance frequency, and f_0 is the linear resonance frequency, determined by fitting a quadratic function to the measured data.

The amplitude frequency coupling in the model (5.5) is equivalent to the normalized nonlinear frequency shift, such that $\alpha = \Delta f / \Delta f_{3\mathrm{dB}}$, where $\Delta f_{3\mathrm{dB}}$ is the resonator's linewidth. This is true because we always take oscillator's uncoupled amplitude to be unity in experiment, effectively scaling out the saturation level, S. The experimentally accessible values of α are plotted as a function of the control parameter in Fig. 5.8. As shown, α can be tuned from 0.11 to 5.93, a range spanning 34.7 dB (~1.5 orders of magnitude). This window of 34.7 dB of dynamic range can be shifted up or down by adjusting the AGC setpoint. The theoretical lowest α is set by the thermomechanical fluctuations of the resonator, and the highest α is set by higher-order nonlinear effects. In Appendix D, we show the resonator can be driven to nonlinear frequency shifts commensurate with $\alpha > 100$ for a duffing nonlinearity, but, at these amplitudes, the model may need to be re-derived with higher-order nonlinear terms.

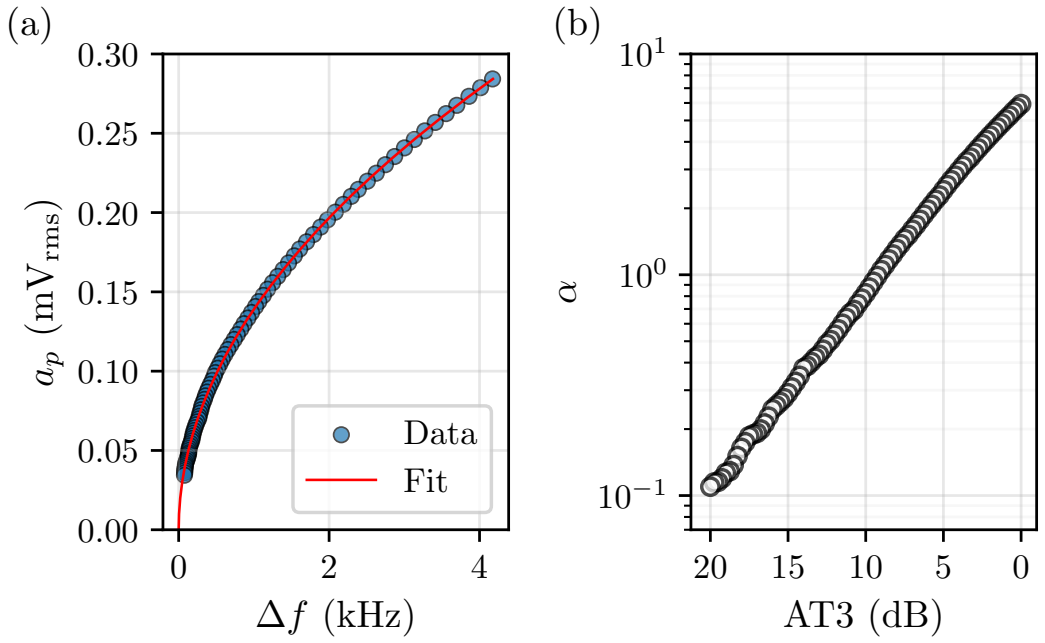

Figure 5.8: Oscillator Nonlinearity Calibration (Experiment). (a) Measurement of the amplitude, a_p, and frequency shift, Δf, of the oscillator as the feedback strength is varied over the electronically accessible range. Each circle represents an experimentally accessible point which is discretized by the possible digital control settings. A quadratic fit is overlaid on the data. $\Delta f = 0$ corresponds to the extracted linear frequency, f_0, from the fit. (b) Logarithmic plot of the effective nonlinearity, α, at each experimentally accessible setpoint shown in (a) versus the logarithmic feedback scaling, set by AT3.

Frequency Fluctuations and Drift

The rich phase noise properties of oscillators in these various feedback configurations were thoroughly studied theoretically in [117] and experimentally in [118]. Additionally, the phase noise reduction occurring from synchronization of two

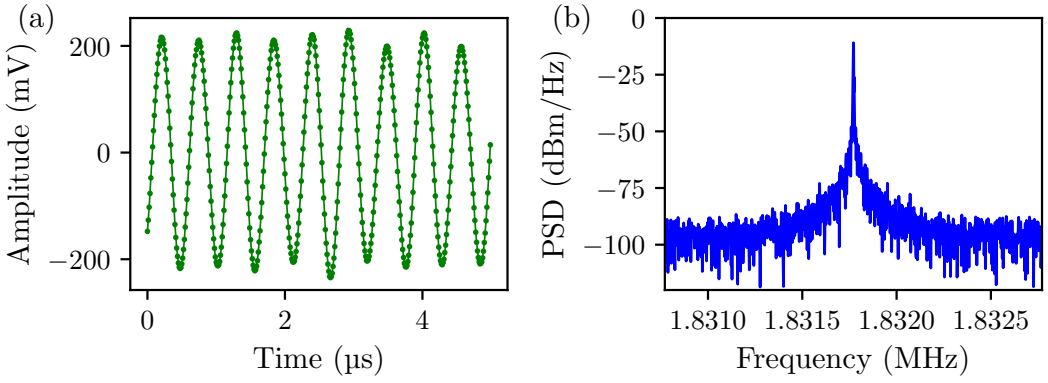

Figure 5.9: Oscillator Characteristics. (a) Dense measurement of the time-domain signal and (b) power spectral density of the oscillation.

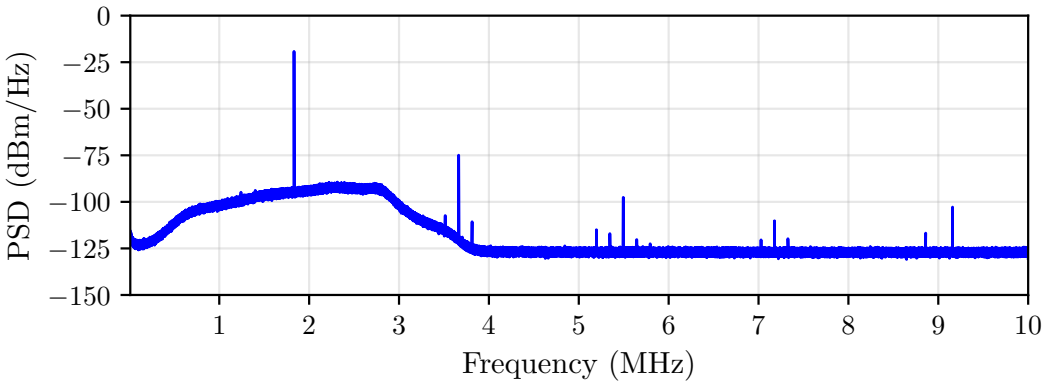

Figure 5.10: Wide PSD. Wide PSD of the oscillator calculated using Welch's method.

nanomechanical oscillators is reported in [119]. In this work, we do not repeat these thorough characterizations, but we measured some of the basic properties of the oscillators for general reference. In Fig. 5.9(a), we show a dense time-domain plot of the oscillator signal, and in Fig. 5.9(b) we show the power spectral density (PSD) in a narrow band around the oscillation frequency. In this plot, the PSD is estimated as the square of the discrete Fourier transform (DFT) components, computed using the fast Fourier transform (FFT) algorithm. This is called a periodogram.

Using the same data, we compute the PSD from 10 kHz to 10 MHz, shown in Fig. 5.10. For this PSD, we use Welch's method to reduce the noise, revealing the higher harmonics of the oscillator, as well as a series of spurious tones. From this PSD, we calculate the 2nd harmonic distortion to be -55.72 dB. Despite being more than two orders of magnitude lower power than the principal tone, the second harmonic may contribute to weak parametric coupling between nodes. To mitigate this, we employ multiple low-pass filters in the coupling paths, but it is possible that there is still some effect, as was described in [28].

The spurious tones come from a charge pump in a component of the phase shifting circuit. In future implementations, we will use a different IC to avoid this; for now, the tones are not near the frequency of the oscillators and do not affect the dynamics. However, these tones and the second harmonic can make it more difficult to cleanly extract the phase of the principal oscillation. Thus, we use both analog and digital filtering of the oscillators to extract the phase dynamics.

Finally, to characterize the frequency fluctuations and amplitude fluctuations, we use a reference phase-locked loop (PLL) to lock to the oscillations and track the oscillator for 4 hours. We set the bandwidth of the PLL to 100 kHz and sample

the amplitude and frequency at 119 samples per second. The normalized amplitude and frequency versus time are plotted in Fig. 5.11(a) and 5.11(b), respectively. The amplitude is plotted as a/a_0 where a_0 is the average amplitude over the interval and a is the instantaneous amplitude. The amplitude is very stable; qualitatively, it fluctuates by $\sim \pm 0.002$ over the entire measurement period, which is a variation of $< 1\,\mathrm{mV_{rms}}$.

The frequency is plotted as the change in frequency from the initial point, Δf, divided by the bandwidth so that the fluctuations can be directly related to the normalized parameters in the theoretical model. We compute Allan deviation of the frequency, σ_f (Hz), divided by the linewidth for various averaging times, plotted in Fig. 5.11(c). This gives a general measure of the long-term parameter drift. For example, over 100 seconds, the frequency drifts by around 0.01 linewidths on average, which is about 7 Hz. Error bars show ± one standard deviation. For more details, see Appendix F.3. The general upward trend in this particular measurement is likely due to slowly drifting ambient room temperature at the time of the measurement.

5.5 Edge Calibration (Experiment)

For each edge, we have to calibrate the phase and the amplitude of the complex coupling coefficient, $|\beta|$. For initial experiments, we aim to have identical coupling across every edge in the network. To do that, we have to set the phase, Δ, the emissivity, κ, and the susceptibility, ν.

We start by setting the phase. Importantly, we cannot control the phase along all 64 possible edges. Thus, we rely on the assumption that the phase error in the switch matrix is small. That is, we assume the real phase difference between electrical paths such as β_{11} and β_{21} is small. In open-loop characterization of the coupling electronics, we find that the phase error at 1.8 MHz is $\sim 0.11°$, but this is measured without oscillators. With oscillators, the worst case input phase error to each is larger but still $<\sim 5°$ (0.087 radians). In our theoretical modeling, we do not account for these random phase lags, but the impact of random or structured phase errors would be an interesting topic for future study.

For a single oscillator, we first set up the switch matrix such that the node is only coupled to itself. Then, we sweep the phase and amplitude of that node's coupling loop path and plot the results as parametric curves in Fig. 5.12. Each colored curve represents a particular coupling strength which corresponds to a particular setpoint

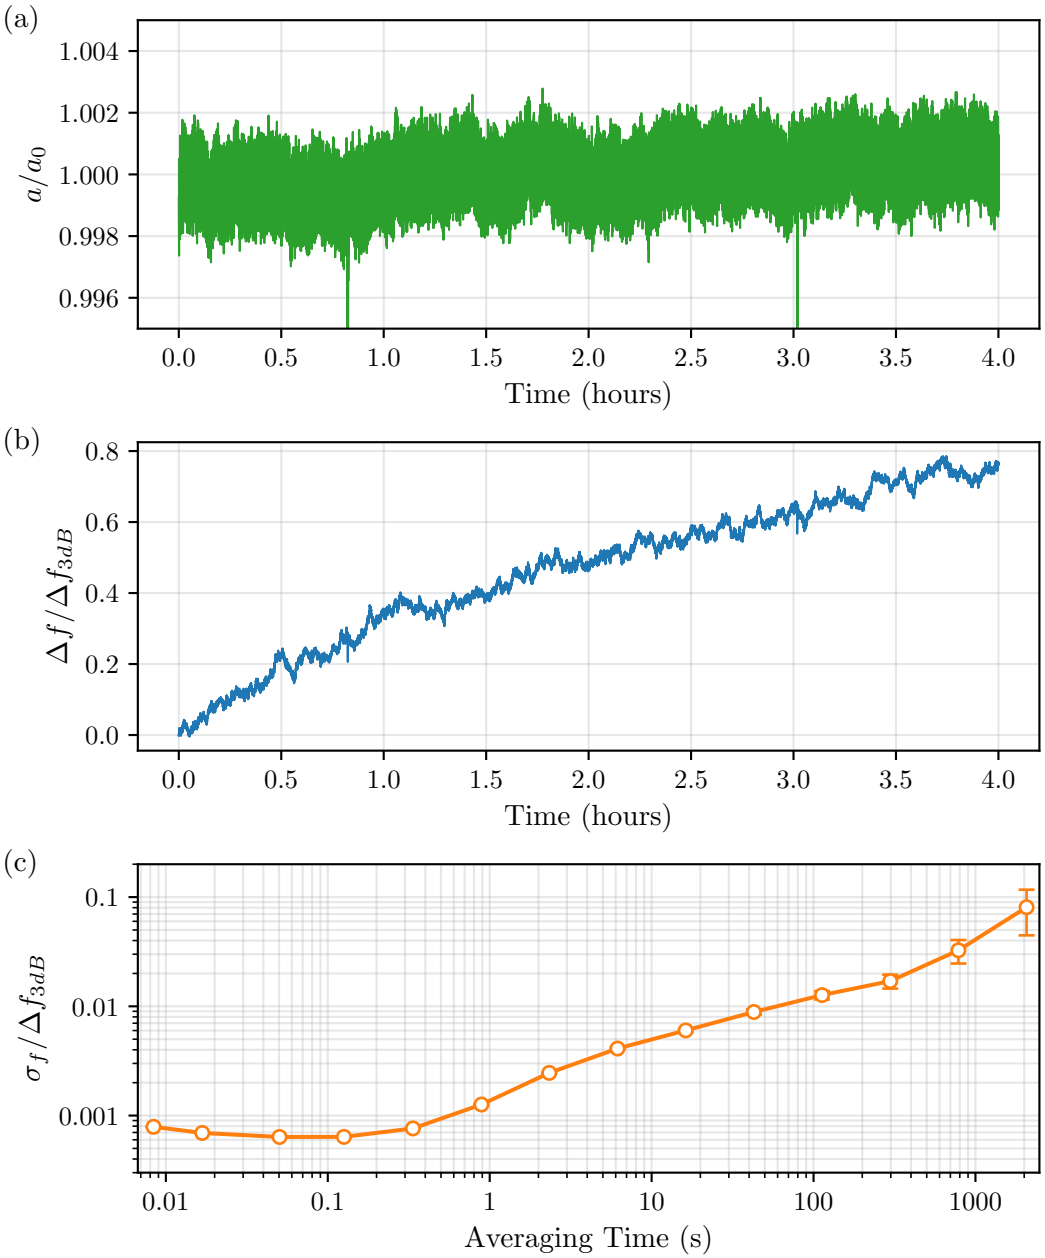

Figure 5.11: Amplitude and Frequency Fluctuations. (a) Normalized amplitude fluctuations and (b) normalized frequency fluctuations as a function of time. (c) Normalized, non-overlapping Allan deviation as a function of averaging time, τ, calculated from (b).

of ATE1. As the phase is varied, we track the coupled amplitude, a_c, normalized to the uncoupled amplitude, a_{uc}, as a_c/a_{uc}, and the normalized frequency shift, $\Delta f/\Delta f_{3dB}$. We see that the oscillator traces out similar curves to those shown in the Fig. 5.5, the theoretical plot of (5.16) and (5.17).

Using this data, we can identify the values of the phase control parameter which most closely match the reactive coupling condition, namely the values where the normalized amplitude is unity. There are two complications: firstly, as previously discussed, the phase shifter is not ideal and also effects the amplitude of the coupling, distorting the curves from the theory. In contrast, the attenuator only has a small effect on the phase. Thus, by varying the phase and attenuation together, we can identify the phase value where the amplitude is invariant to the coupling attenuation. This signifies that the coupling is purely reactive, resulting in only a frequency shift regardless of the value of $|\beta|$. The second complication is noise and drift. The oscillators are not temperature stabilized, so the frequency drifts noticeably on the timescale of seconds. We can compensate for this by constantly measuring the uncoupled frequency to correct for drift. Additionally, the amplitude is very stable and allows for accurate calibration of the phase.

To calibrate the amplitude of the coupling path, we set the phase to the reactive coupling point. Then, we sweep the attenuator ATE1 and measure the normalized frequency shift. From Table 5.1, the frequency under reactive-attractive coupling is $\Omega = \alpha - |\beta|/2$. Therefore, using the frequency shift with the coupling on versus off, $|\beta| = 2\Delta\Omega = 2\Delta f/\Delta f_{3dB}$. Using this relationship, $|\beta|$ is plotted as a function of the edge attenuation, ATE1, in Fig. 5.13. For clarity, the attenuation is incremented by 0.5 dB in this plot, but the step resolution of the experiment is 0.25 dB. So, the number of achievable states is double what is shown.

In this plot, we measure the frequency shift at each point 15 times. The error bars show one standard deviation in $|\beta|$ from those repeated measurements, per side. At large values of $|\beta|$, the error is small relative to the frequency change. But, at low values of $|\beta|$, the frequency shifts are small, meaning that frequency fluctuations and drift over the measurement timescales introduce significant errors. Additionally, we use a discrete Fourier transform to extract the frequency. At low values of $|\beta|$, the DFT binning can introduce quantization error. This creates a tradeoff between the amount of drift between measurements and the frequency resolution. Fast measurements result in less drift but more bin error, and slow measurements result in more drift but less bin error. This effect can be mitigated with various tactics. For

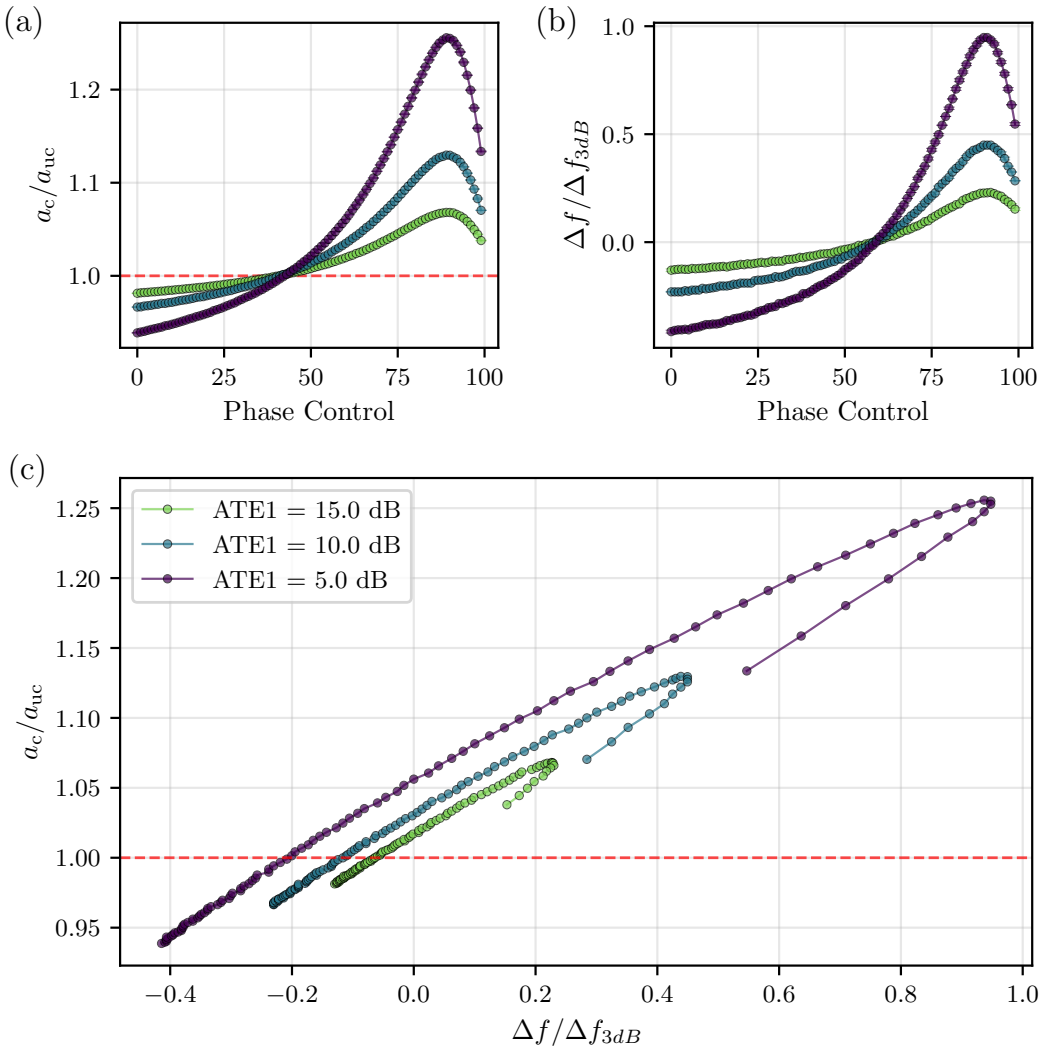

Figure 5.12: Parametric Sweep of Edge Phase (Experiment). Plot of oscillator's amplitude versus the oscillator's frequency shift as a function of the coupling phase and amplitude for a fixed nonlinearity, $\alpha = 1.69$. Each curve is a sweep of experimentally accessible coupling phase for a given setpoint of the attenuator, ATE1. We only show a small subset of ATE1 setpoints for clarity. On the y-axis, the amplitude is plotted as a ratio of the amplitude when the coupling loop is off, a_{uc}, to the amplitude when the coupling loop is on, a_c. The frequency is plotted as the normalized frequency shift, $\Delta f/\Delta f_{3\text{dB}}$, from when the coupling loop is on to when it is off, divided by the resonator's linewidth. A red dashed line indicates where $a_c/a_{uc} = 1$, indicating that the effect of the coupling loop is a pure frequency shift and the coupling phase is purely reactive. These curves correspond to the experimentally accessible portion of the theoretical curves shown in Fig. 5.5 for a few select β values.

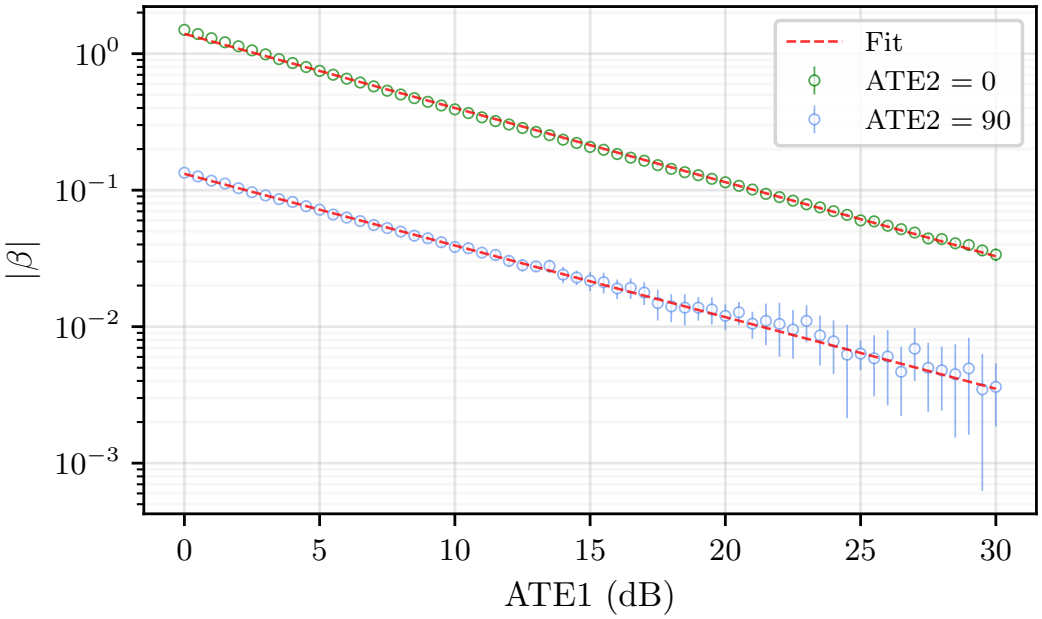

Figure 5.13: $|\beta|$ Calibration (Experiment). Plot of the coupling magnitude, $|\beta|$ of the edge as a function of the first edge attenuator setting, ATE1, for two values of the secondary edge attenuator, ATE2. For both, the phase is fixed at the value corresponding to the reactive coupling point. Here, ATE1 is plotted in steps of 0.5 dB, but the experiment allows for steps as small as 0.25 dB. Error bars represent the standard deviation of the frequency shift measurements over 15 runs. For each value of ATE2, a red line shows an exponential fit which is used as a calibration curve in experiments. For the data corresponding to ATE2=90, only points with ATE1 < 15 dB are used to fit the calibration curve and remaining values of ATE1 are extrapolated.

example, we cancel long time drifts by measuring the uncoupled frequency before each shift measurement, and we fit the DFT peaks with quadratic interpolation to obtain approximate sub-bin resolution without extending the measurement time. Furthermore, when operating in the low $|\beta|$ regime (e.g., ATE2=90), we find that it is more reliable to extract a fit of the $|\beta|$ curve using the low error values from ATE1=0 to ATE1=60. In-situ characterization measurements of the attenuator show strong linearity and low attenuation error across the operating regime. Thus, the low $|\beta|$ measurements are accurately predicted by the fit.

For future experiments, it could be beneficial to utilize external frequency tracking tools, for example by locking a reference PLL to the NEMS oscillator output, to quickly track the frequency. This could substantially reduce calibration time and increase accuracy. The advantage of the method employed here is that it is readily

achieved with only the data logging tool — the oscilloscope — and digital signal processing.

By varying ATE1 and ATE2, the system can achieve a range of coupling strengths from approximately 0.004 to 1.495, a 51.5 dB range. Across the entire network, this dynamic range is slightly reduced because of variations across the network edges, and a typical range of $|\beta|$ achievable across all edges is 0.005 to 1.2, a 47.6 dB range. This range is purely a function of the electronics and can be changed by designing different circuits or modifying the passive components on the existing PCB.

There is one additional calibration step to ensure that all nodes have matched susceptibility, which is controlled by ATE2. Small variations in electrical impedance and device properties lead to mismatches on the level of 1 dB. To correct for these, we configure the system in a one-to-all configuration and directly measure how all other nodes respond to injection locking from a single node. This directly measures variations in susceptibility across the nodes which we then correct for by varying ATE2.

To turn the coupling off, it is insufficient to maximize ATE1 and ATE2. Instead, the off state is achieved by turning off the switch matrix which reduces the coupling to parasitic electrical effects. These parasitic effects set the true minimum $|\beta|$. The electrical crosstalk between channels of the network card was measured to be -75 dB at 1.8 MHz, but this measurement is done without the oscillator PCBs. We do not further characterize the crosstalk, but it is sufficiently low that the oscillators do not stably synchronize when the coupling is turned off, even when they are nearly identical in frequency. However, the crosstalk may still affect the phase evolution, resulting in a bias in the initial conditions of experiments.

5.6 Conclusion

In this chapter, we carefully reviewed both the theory and practice of how the network parameters are calibrated. We showed how the effective nonlinearity, α, can be tuned from 0.11 to 5.93, and how the coupling strength, β, can be made reactive or dissipative and tuned over a range from 0.004 to 1.495. These methods are frequently employed in the experimental system to maintain the calibration and instantiate novel dynamics.

Chapter 6

SYNCHRONIZED AND UNSYNCHRONIZED STATES OF UNIFORM COMPLEX GRAPHS

In this chapter, we study the steady-state dynamics of the systems on various complex graphs and in different regions of the parameter space. Namely, we explore circulant graphs, wheel graphs, and Erdős–Rényi random graphs with uniform parameters and identical frequencies. Primarily, we study the case of $\alpha = 2$ with $0.04 < \beta < 1.0$. In the experiment and in simulation of the full amplitude-phase model, we found numerous stationary and nonstationary states which vary widely depending on the underlying network topology.

6.1 Experiment Protocol

In this chapter, we study the dynamics of the system as it evolves from an uncoupled state to a coupled state under a static graph. To achieve this, we first calibrate all of the parameters of the system, as described in Chapter 5 and program the attenuators and phase shifters to the desired state. The switch matrix is set to off, effectively uncoupling all of the nodes. The uncoupled nodes drift in frequency. So, before turning the coupling on, we measure the frequency of each node and apply a correction to their individual DC biases. This brings them into frequency alignment with a maximum spread < 0.01 linewidths.

Then, the switch matrix is abruptly turned on, coupling the nodes according to the programmed connectivity. The network then evolves from the initial state to one of many possible final states, some of which are stationary and some of which are time varying. The initial phases are quasi-random due to the slow frequency drift and random start time; however, we observe some small correlations between nodes due to the parasitic coupling. This effectively biases the initial conditions, leading to a biased sampling of the final states.

We start the data collection shortly before the coupling is turned on and record the full nodal dynamics over thousands slow-time units, which corresponds to one or two seconds and millions of samples. From this data, we extract the phase and amplitude of the primary oscillation using a Hilbert transform. We repeat this procedure N_R times per graph to observe the attractors with the largest basins of

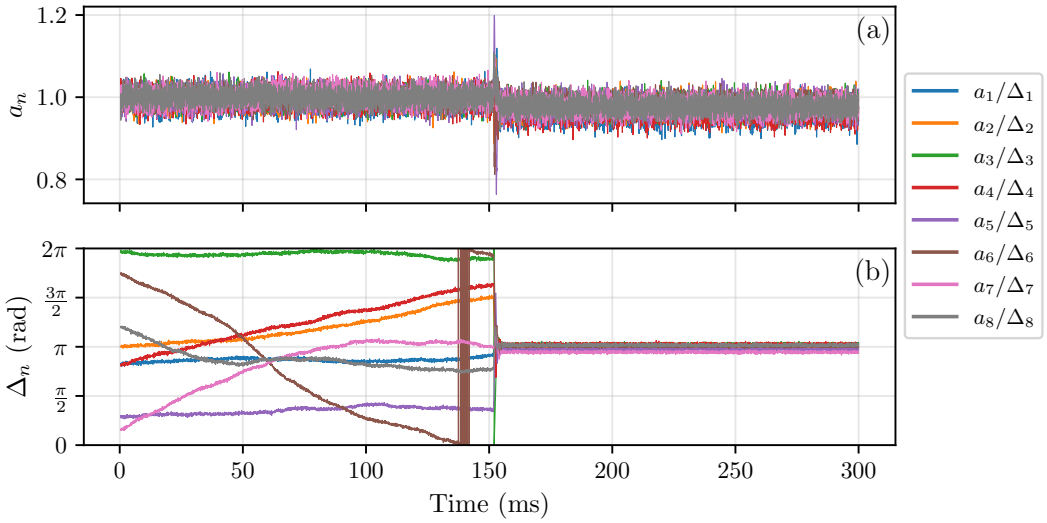

Figure 6.1: Time Domain Synchronization on a Ring. Example time-domain data of the (a) normalized amplitude and (b) phase differences from a ring of 8 oscillators. Before the coupling is turned on, the phases drift relative to each other according to the frequency mismatch between oscillators. Just before t~151.8 ms, the coupling is turned on and the oscillators dynamically synchronize, in this case to the anti-phase state. Note that the phase is plotted modulo 2π, so there can be artificial jumps where an oscillator crosses the boundary.

attraction. Typically, $N_R = 200$, but for some parameters sweeps we use $N_R = 100$. The value will be indicated in the caption of the figure. From this method, we do not expect to observe all possible states as some, such as inhomogeneous states, are unlikely to arise from random initial conditions [28].

In the case of circulant graphs, it is sensible and convenient to consider the phase difference between nearest-neighbors, so we typically plot the data using $\Delta\phi_n = \phi_n - \phi_{n-1}$. For the sake of brevity, we do not show all of the data for all graphs, we only show a selection of the most relevant data. Fig. 6.1 shows an example of the time-domain data collected from an experiment conducted on a ring. From $t = 0$ and $t = 151.7$ ms the coupling is off and the phase differences drift in time due to small frequency differences. At $t = 151.7$ ms, the coupling is abruptly turned on. The network goes through a transient period before relaxing to the anti-phase synchronized state where $\Delta_i \approx \pi$ for all i.

6.2 Circulant Graphs

We use the network to study the dynamics of connected circulant graphs with reactive-repulsive coupling. For initial experiments, we set $\alpha = 2$ and $\beta = 0.2$, which we expect will result in dynamics generally represented by the first-order

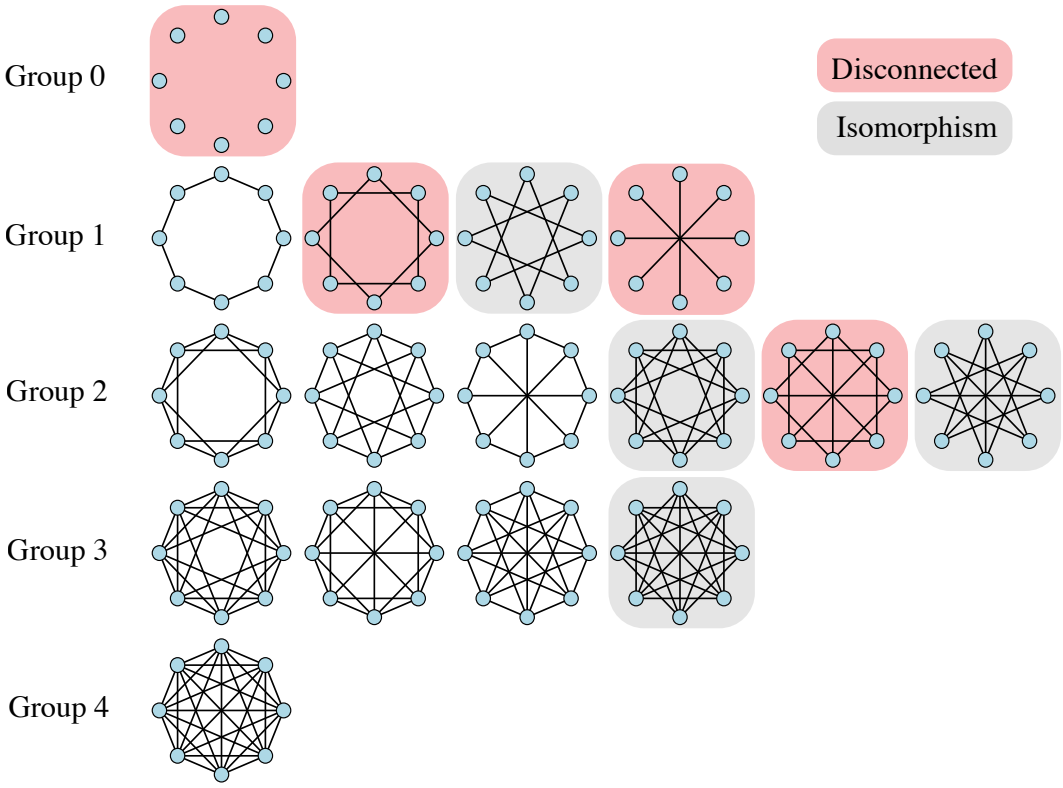

Figure 6.2: 8-Node Circulant Graphs. All circulant graphs with N=8, organized by the number of neighbor connections.

phase model. Here, we primarily focus on the experimental results, but numerical simulations were carried out for each graph to verify the results.

For 8 nodes graphs, it is possible to draw 16 circulant graph representations. In Fig. 6.2, we shown each of these representations organized by the number of neighbors they are connected to, which we call their group. We denote the circulant graphs as $C_N(S)$ where $N = 8$ is the number of nodes and S is the set of connections, $S \subset \{1, 2, \ldots, \lfloor N/2 \rfloor\}$. For example, $C_8(\{1, 2\})$ denotes a circulant graph with nearest-neighbor and next-nearest-neighbor connections. All circulant graphs are undirected.

First, group 0 has the null graph, which is not connected so we discard it. At degree one, there are 4 graph representations, but $C_8(\{2\})$ and $C_8(\{4\})$ are unconnected and $C_8(\{3\})$ is isomorphic to $C_8(\{1\})$. Thus, only $C_8(\{1\})$ is of interest. Group 2 has 6 graphs. $C_8(\{2, 4\})$ is not connected. $C_8(\{2, 3\})$ is isomorphic to $C_8(\{1, 2\})$, and $C_8(\{3, 4\})$ is isomorphic to $C_8(\{1, 4\})$. Therefore, we only need to consider the dynamics of $C_8(\{1, 2\})$, $C_8(\{1, 3\})$, and $C_8(\{1, 4\})$. In group 3, $C_8(\{2, 3, 4\})$ is isomorphic to $C_8(\{1, 2, 4\})$, leaving $C_8(\{1, 2, 3\})$, $C_8(\{1, 2, 4\})$, and $C_8(\{1, 3, 4\})$.

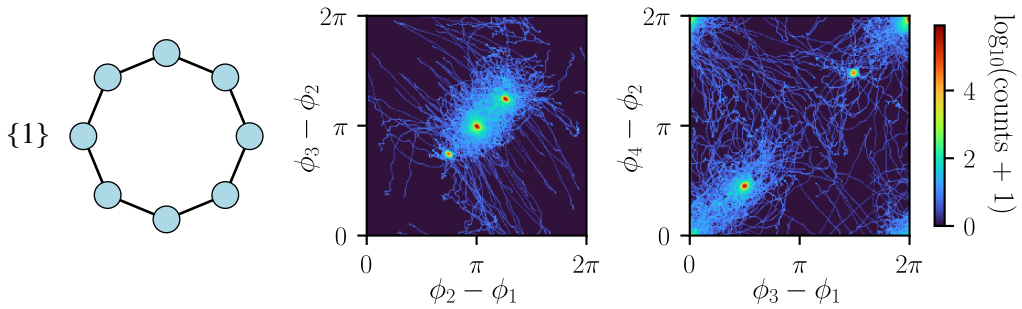

Figure 6.3: Synchronized States of a Ring (Group 1). (left) the set of edges followed by the graph, (middle) heatmap showing a NN slice of phase space where $N_R = 200$ system trajectories evolve from random initial conditions to one of three synchronized fixed points. (right) NNN slice of phase space showing the same. The color indicates the logarithm of the density of points at that point in phase space. Here, $\alpha = 2.0, \beta = 0.2$ and $N_R = 100$.

Finally, group 4 only has $C_8(\{1, 2, 3, 4\})$ which is the complete graph. We include it in our study of circulant graphs, but it is a special case because it not only obeys the D_8 symmetry of the other graphs, but also exhibits full permutation symmetry described by the symmetric group, S_8.

All circulant graphs have a spatial symmetry described by the dihedral group, D_N.

In group 1, there is only the ring graph. The results of the experiment are shown in Fig. 6.3. This figure shows a heatmap which is composed of all the experimental runs; in this case, N_R=200. From that data, we can see that there are three fixed points which make up the largest basins of attraction for the system. These stationary points are the splay states and were thoroughly explored in [28]. In this high α, low β regime, we only observed these stationary states in both experiment and simulation.

In group 2, all of the graph representations can fully represented by three unique graphs, and the experimental data for these graphs is shown in Fig. 6.4. Here, we carry out the same experiment with $\alpha = 2.0$, $\beta = 0.2$, and $N_R = 200$ for each graph. For $C_8(\{1, 2\})$, we find more than one possible state. First, in the experiment, we observe one fixed point where nearest neighbors have a fixed phase shift of $\Delta_i = 5\pi/4$. In simulation another splay state is observed at $\Delta_i = 3\pi/4$. Likely, this splay state would be observed in experiment if more runs were conducted. However, the anti-phase synchronized state is not observed in experiment or simulation for this graph, indicating that the NNN connections tend to shrink the basin of attraction or make the state unstable entirely.

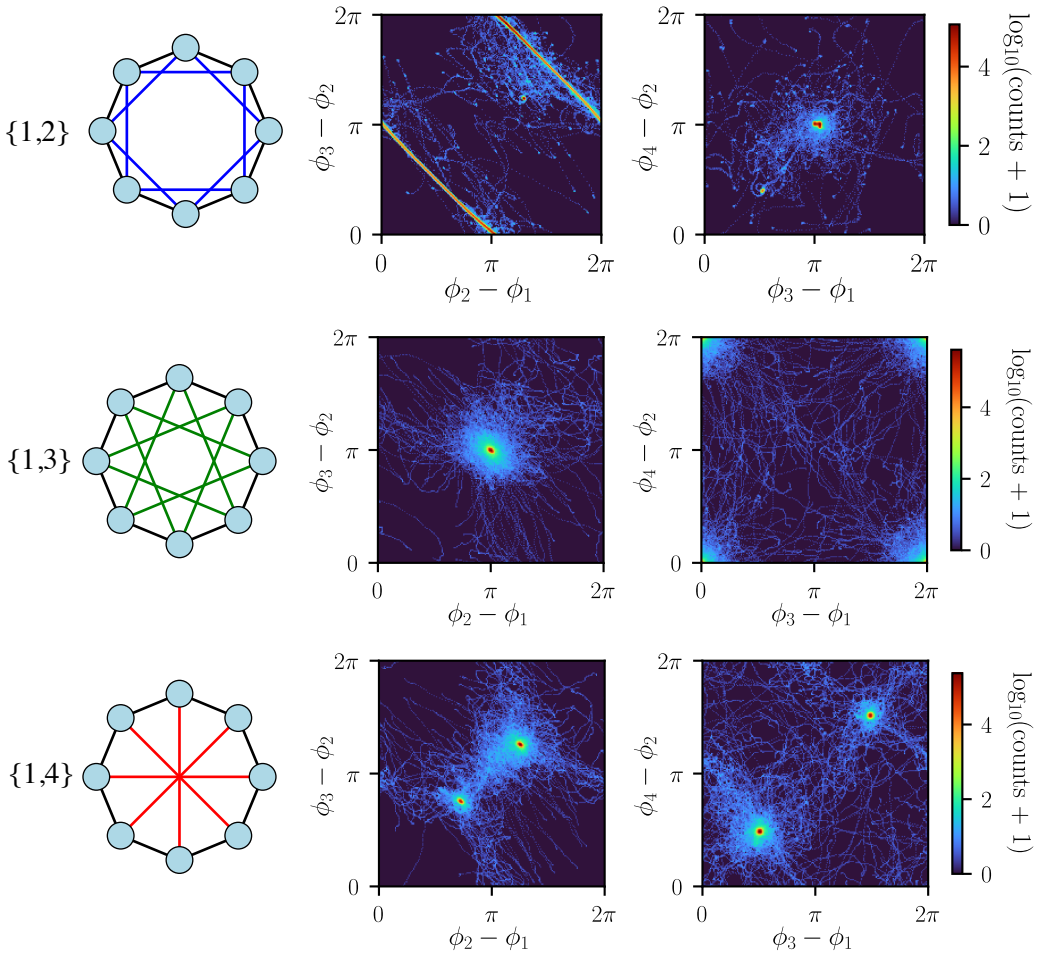

Figure 6.4: States of Second-Order Circulant Graphs (Group 2). (left) the edge set and the graph, (middle) NN phase space slice, and (right) NNN phase space slice. For C_{12}, $N_R = 100$, but for C_{13} and C_{14}, $N_R = 200$. For all data presented here, $\alpha = 2$ and $\beta = 0.2$.

In addition to the fixed points, there is also a decoupled state, which appears as a diagonal streak across the NN phase plane. This state was also observed in [28] and has been explored theoretically in [4]. In Fig. 6.5(a) we show the time-domain evolution of one of these states. After the coupling is initiated, the system goes through a transient before splitting into two distinct groups. By examining the instantaneous frequencies of the oscillators in Fig. 6.5(b), we observe that there are two clusters composed of $n \in \{1, 3, 5, 7\}$ and the other of $n \in \{2, 4, 6, 8\}$. Here, $\Delta f = f_n - \bar{f}$ indicates the frequency relative to the average frequency of all the oscillators, $\bar{f}$, in Hz.

By examining the mutual information in Fig. 6.5(c), we can see that the clusters are generally uncorrelated. The mutual information between NNN is much higher

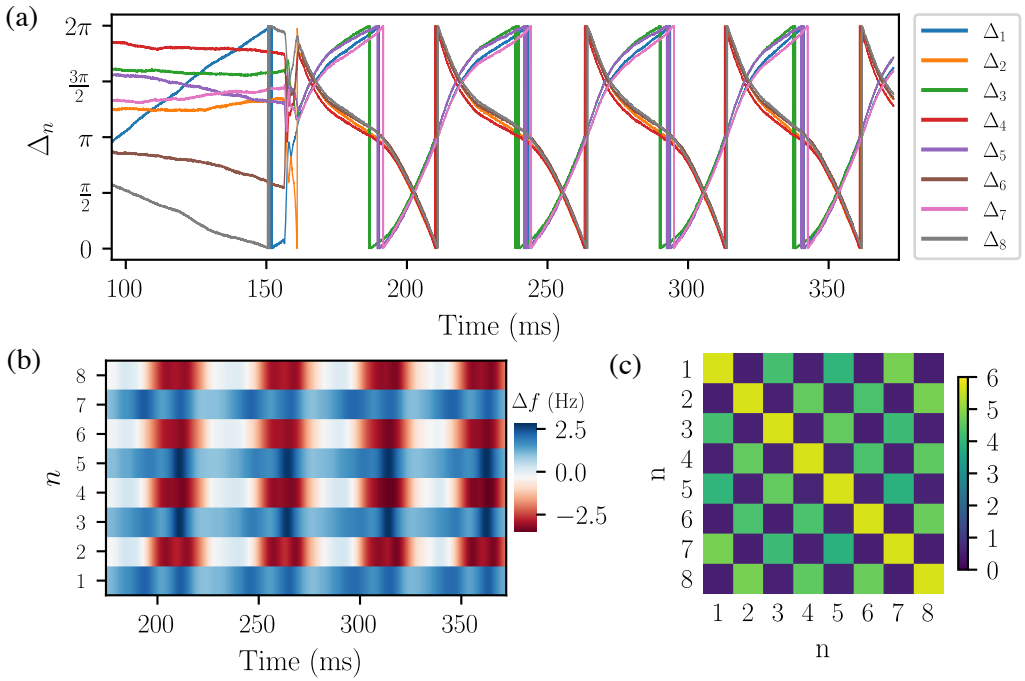

Figure 6.5: $C_8(\{1,2\})$ **Decoupled State Details.** Details of the decoupled state observed for the $C_8(\{1,2\})$ network. (a) time-domain evolution of the NN phase differences in the decoupled state. (b) frequency-spacetime plot showing the frequency evolution vs. time. (c) mutual information matrix.

than that between NN. However, the mutual information between NN is not zero, indicating that the clusters are not perfectly decoupled. This can also be seen from Fig. 6.5(a) and (b) where the phase differences and frequencies do not evolve uniformly. Instead, the clusters phase-slip against each other while maintaining an average frequency difference of around 5 Hz. In simulation, the clusters are perfectly decoupled if the network parameters are perfectly uniform. Small variations to the symmetry of the coupling prevent the clusters from fully decoupling. Further, as reported in [28], the frequency difference between the clusters is a function of the intrinsic frequencies of their constituent oscillators. Therefore, for larger variations in frequency, the clusters precess around each other faster. For perfectly uniform frequency, simulations indicate that the clusters do not precess at all.

While this state was observed in experiment previously [28], here we have realized it using physical NNN connections whereas previous experiments relied on emergent higher-order interactions in the low α, high β parameter regime. Here, we have validated that such a state is possible to realize in the high α, low β regime by explicit construction of the higher-order interactions.

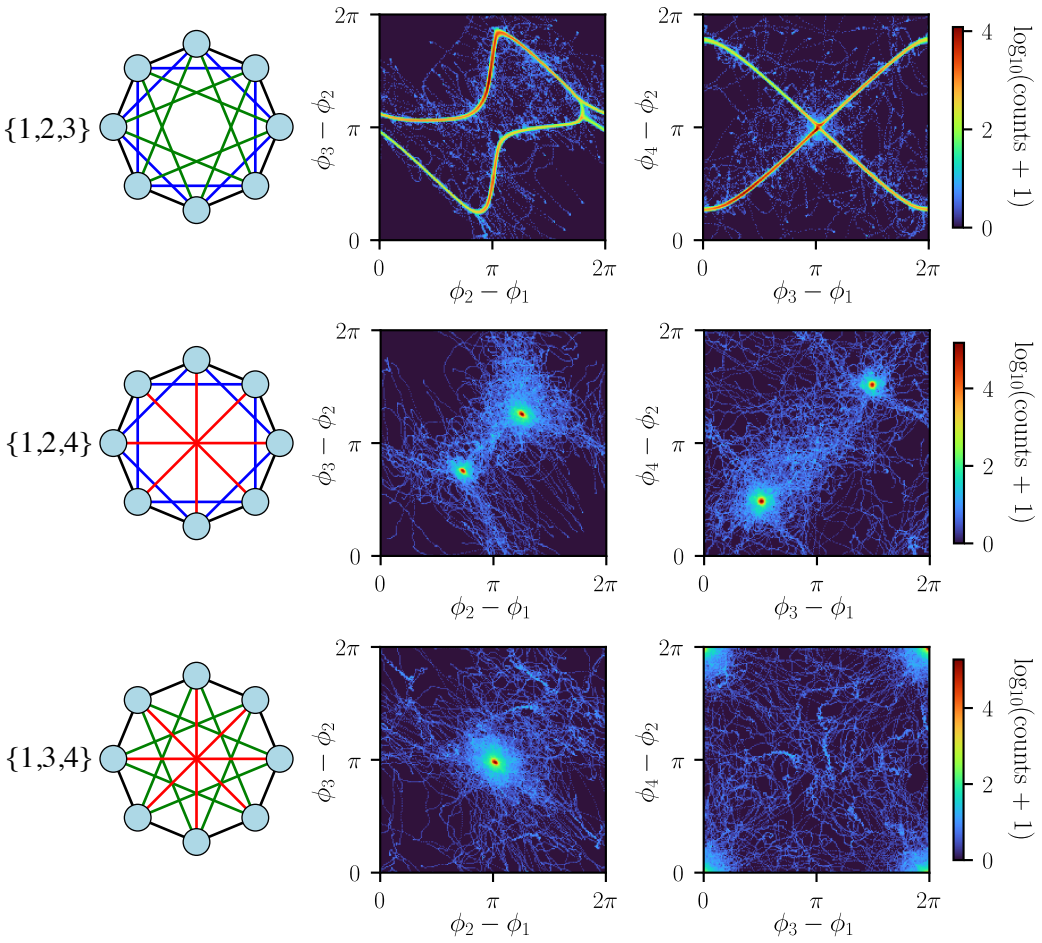

Figure 6.6: States of Third-Order Circulant Graphs (Group 3). (left) the edge set and the graph, (middle) NN phase space slice, and (right) NNN phase space slice. For C_{123}, $N_R = 100$, but for C_{124} and C_{134}, $N_R = 200$. For all data presented here, $\alpha = 2$ and $\beta = 0.2$.

The other two graphs in group 2, $C_8(\{1,3\})$ and $C_8(\{1,4\})$ produce only fixed points states in both the experiments (Fig. 6.4) and simulation. However, on $C_8(\{1,3\})$ only the anti-phase synchronized state appears to be stable. In contrast, on $C_8(\{1,4\})$ the $k = 3$ and $k = 5$ splay states are both stable and the anti-phase state is not observed.

In group 3, there are three unique graphs which are represented by C_{123}, C_{124}, and C_{134}. The experimental state space trajectories for these graphs are shown in Fig. 6.6. Starting with C_{123}, we observed no fixed point solutions in experiment or simulation. We observed only a nonstationary state where the oscillators evolve in 4 clusters, each composed of a set of opposite pairs. In Fig. 6.7(a) the time-domain evolution of the phase differences is shown. Two of the clusters are frequency

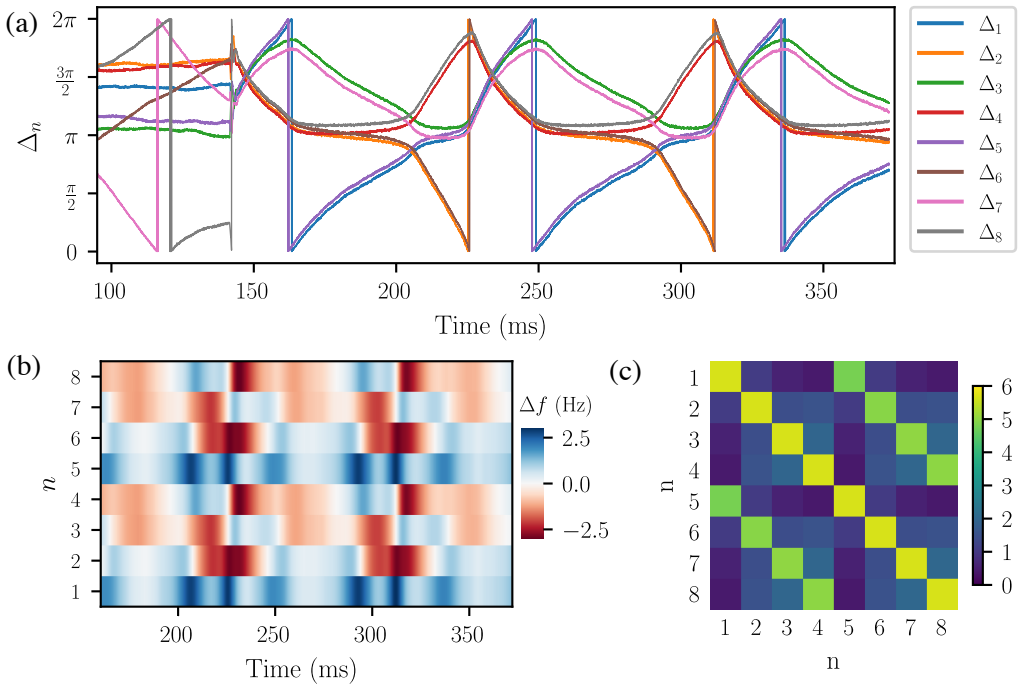

Figure 6.7: C_{123} **State Details.** Details of the complex state observed on the C_{123} graph. (a) time-domain evolution of the phase differences, (b) frequency-spacetime plot, and (c) mutual information matrix.

synchronized but not phase synchronized. The other two clusters are not frequency synchronized to each other or to the other clusters. One has a higher frequency than the main cluster and one is lower. From the frequency-spacetime plot (Fig. 6.7(b)) and the mutual information matrix (Fig. 6.7(c)), the clustering between opposite pairs is apparent. Simulations corroborate the general structure of this state, but more work is needed to understand the dynamics and symmetries.

For the other two graphs in group 3, C_{124} and C_{134}, we only observe fixed point states in both experiment and simulation. In C_{124} we observe two fixed points, the $k = 3$ and $k = 5$ splay states. On the other hand, for C_{134}, we only observe the anti-phase synchronized state. These results are corroborated by simulation of these networks.

Finally, in Fig. 6.8, we show the heatmap for the complete graph under the conditions of reactive-repulsive coupling with $\alpha = 2.0$ and $\beta = 0.2$. Here, we find that none of the oscillators synchronize, which is indicated by the lack of order in the heatmap of the trajectories.

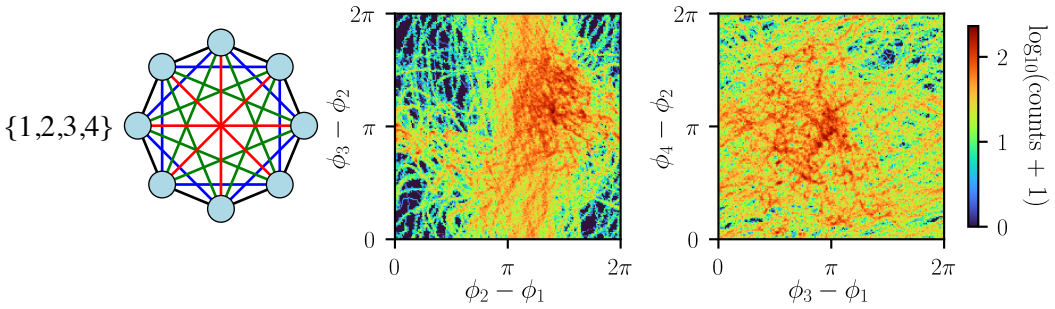

Figure 6.8: Complete Graph States (Group 4). (left) connection indices and graph of the complete graph. (middle) NN phase difference heat map, and (right) NNN phase difference heat map.

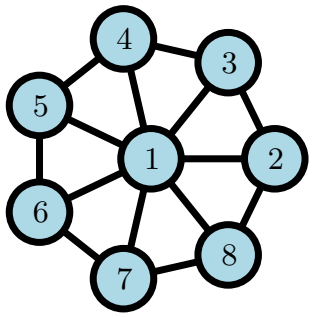

Figure 6.9: Wheel Graph. Wheel graph with 7 outer nodes and 1 hub.

6.3 Wheel Graph

The wheel graph is shown in Fig. 6.9. It is composed of 8 nodes, where the 7 outer nodes make up a ring graph in which each node is also connected to a central hub node. In this case, the center node has the index 1, and the outer nodes are labeled 2 through 8. This graph has the spatial symmetry of the D_7 group.

For this graph, we perform experiments and simulations of the transition to synchronization under reactive-repulsive coupling with $\alpha = 2.0$ and $\beta = 0.2$. In Fig. 6.10, we show an example of the time-domain evolution around the time that the coupling is switched on. We see that there is an initial transient period followed by a slow collapse to a synchronized state. In Fig. 6.10(b) we show the data for 200 runs starting from random initial conditions. We can see that there is a clear structure in the route to synchronization and that the system tends to follow a certain attractor before stopping at a point along the trajectory.

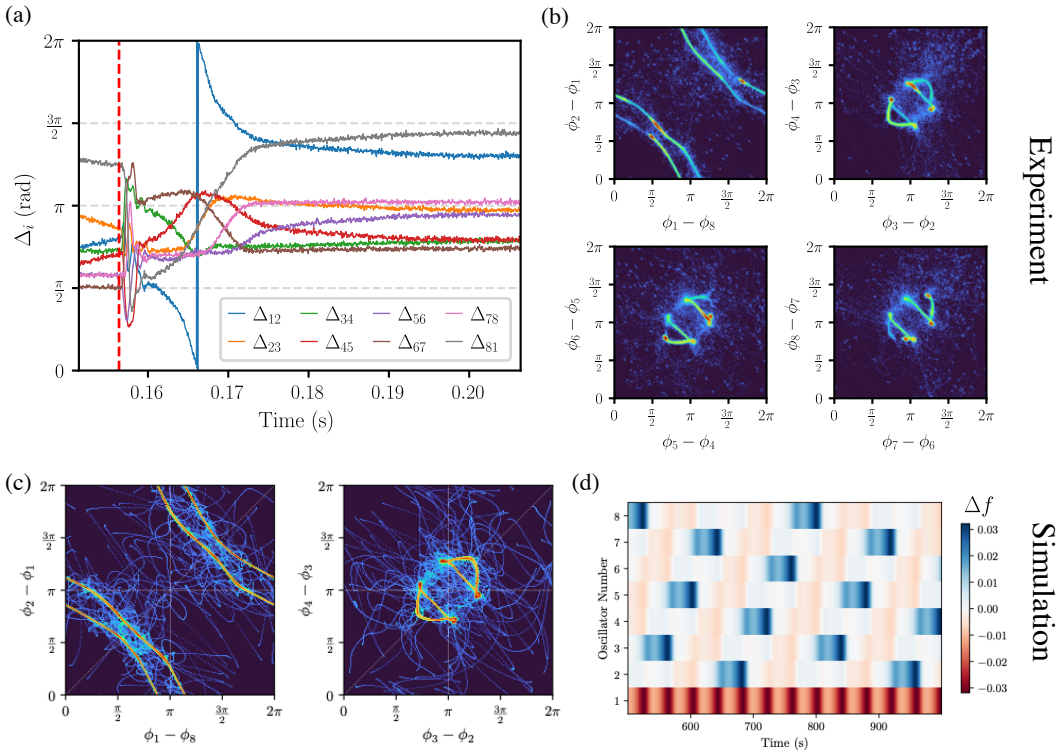

Figure 6.10: 8-Node Wheel Graph States. (a) Example of time-domain evolution of the network when the topology is switched from uncoupled to a hub graph network. (b) Phase space evolution of 200 random initial conditions. (c) simulation of the full amplitude-phase model for 200 random initial conditions. (d) Frequency spacetime plot from simulation of network after collapsing to the traveling wave state.

To investigate this, we simulated the full amplitude-phase model with no frequency dispersion. The resulting trajectories are shown in Fig. 6.10(c). Here, we see the same structure, but the oscillators do not converge to a fixed point. Instead, their phases cycle indefinitely around in a traveling wave pattern. For one of the simulation runs, we plot the frequency spacetime data which shows that oscillator 1 is not frequency synchronized to the others but does have some modulation. The other 7 show a traveling wave pattern which traverses the ring twice before repeating.

In the experiment, we do not observe continued cycling of the traveling wave. This could be due to small parameter or frequency variations which cause the system to get stuck in a fixed point appearing along the wave's trajectory. Simulations indicated that a frequency spread of as little as 0.01 linewidths can cause the wave to stop. Further investigation of this state would be interesting to examine its exact characteristics and sensitivity to real-world conditions. In this parameter regime, this traveling wave state does not exist on the ring graph alone. That suggests that

the center node creates an instability or frustration which leads to the pulse traveling around the outer ring. Additionally, the hub node may be a source of long range interactions which facilitate the existence of the state.

6.4 Erdős–Rényi Graphs

As introduced in §3.3, we can leverage randomness and criticality to construct reservoir computers out of coupled oscillator networks. Towards that goal, we conducted experiments on undirected Erdős–Rényi (ER) graphs. An ER graph is a type of random graph which is constructed by considering every possible edge in a network and including the edge with an independent probability, $p \in (0, 1)$. Although with only 8 nodes the graphs are subject to large statistical fluctuations. We studied their synchronization behavior with both reactive-repulsive and reactive-attractive coupling. We found that multiple synchronized states are possible with the repulsive case, but the attractive case usually leads to the in-phase synchronized state.

Example 1: Repulsive Coupling

Fig. 6.11 shows the results from a specific ER graph constructed with $p = 0.7$, reactive-repulsive coupling, and $\alpha = 2$, $\beta = 0.2$. Over $N_R = 100$ experimental runs, we observe only two states, both of which are fixed points of the dynamics. The heatmaps show an arbitrarily chosen partial slice of the phase space. While not fully indicated here, we verify that the system only converges to one of two possible fixed points, one with a probability of 66%, shown in Fig. 6.11(c) and the other with a probability of 34%, shown in 6.11(d). Assuming a binomial distribution, the 95% confidence interval for the probability of stabilizing in state (c) is given by $p[(c)] \in [0.558, 0.752]$. This indicates a likely bias towards that state.

Beyond the steady-state solution, heatmaps show the convergence to the fixed points predominantly along certain directions in the phase space, indicating the presence of attracting manifolds along which the system prefers to evolve. Additionally, there appear to be trajectories which cross between the fixed points. These are likely cases where part of the network begins or evolves near one fixed point while other nodes are far away. However, as the other nodes evolve they pull the entire network into the alternative fixed point and thus part of the system leaps across, leaving a streak in the probability density. However, this phenomenon requires further study.

In experiment, we sweep β from 0.01 to 0.4, and across that range, no qualitative changes to the state space structure are observed. However, we observe a decrease

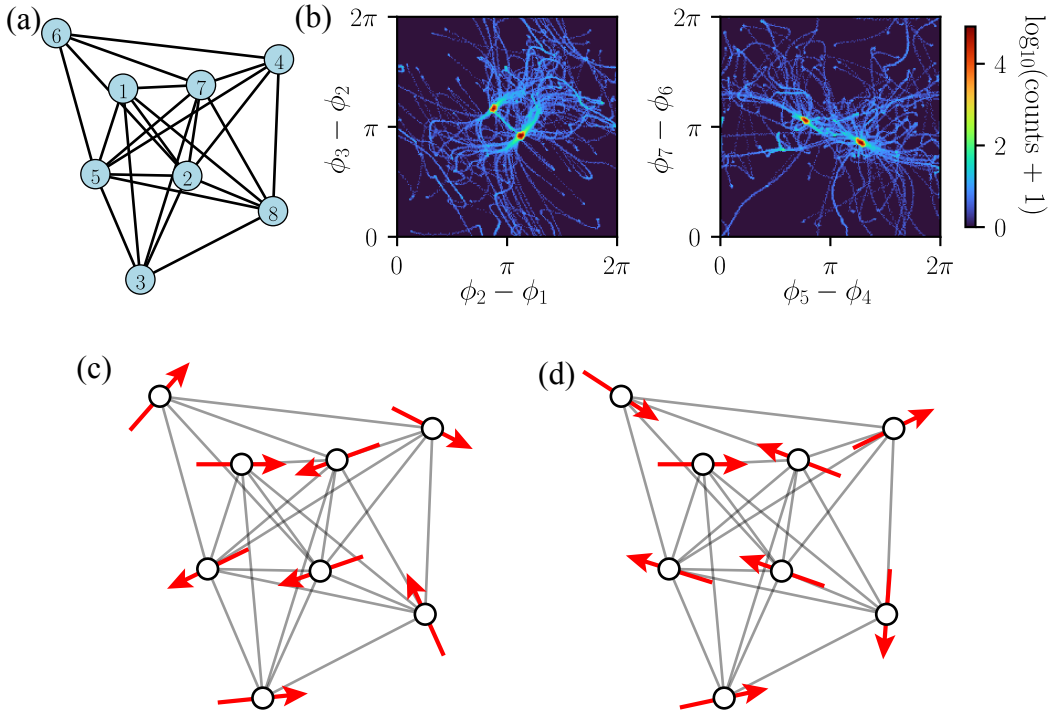

Figure 6.11: Synchronized States of an ER Graph with Reactive-Repulsive Coupling. Experimental data showing the synchronization transition on an Erdős–Rényi random graph with $\alpha = 2$ and $\beta = 0.05$. (a) schematic of the undirected random graph generated by adding each edge with a probability of $p = 0.7$. (b) Two slices of phase space, showing $N_R = 100$ experimental trajectories starting from random initial conditions, shown as a heatmap. (c) and (d) diagrams showing the spatial phase configurations of the two observed fixed points after the system reaches steady-state. In each, the phase of oscillator 1 is set to $\phi_1 = 0$ and the other phases are shown relative to ϕ_1.

Table 6.1: Counts of State (c) vs. State (d) for Varying β (Experimental Data).

Exp. #	β	(c)	(d)	$C_{0.95}(p_c)$
1	0.01	60	40	[0.50, 0.70]
2	0.05	66	34	[0.56, 0.75]
3	0.1	50	50	[0.40, 0.60]
4	0.2	47	53	[0.37, 0.57]
5	0.3	49	51	[0.39, 0.59]
6	0.4	55	45	[0.45, 0.65]

in the time to synchronization and a larger frequency shift of the synchronized state. Detailed analysis of these effects will be the topic of future study. In Table 6.1, we report the statistics collected for running the experiment at 6 different values of β.

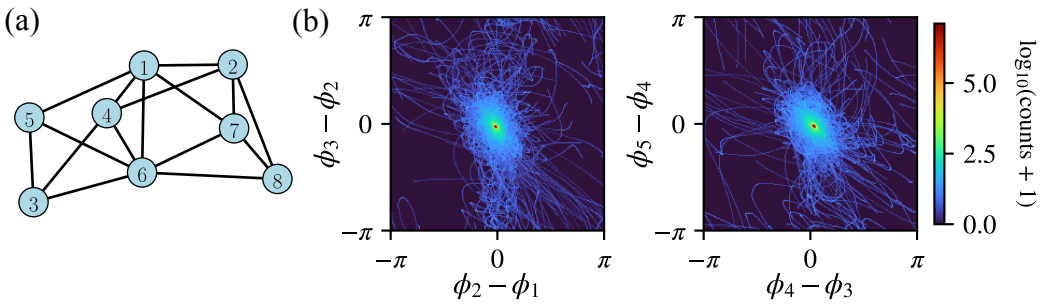

Figure 6.12: Synchronized States of an ER Graph with Reactive-Attractive Coupling. Experimental data showing the synchronization transition on an Erdős–Rényi random graph with $\alpha = 2$ and $\beta = 0.2$. (a) schematic of the undirected random graph generated by adding each edge with a probability of $p = 0.5$. (b) Two slices of phase space, showing $N_R = 200$ experimental trajectories starting from random initial conditions, shown as a heatmap.

To complement the experiments, we simulated the full amplitude-phase model with this graph and with matching nonlinearity and coupling. Qualitatively, the results match well. The same two fixed points are observed in the simulations with similar structure to the trajectories. For $\beta = 0.05$ and $\alpha = 2$, we simulated the network from 1000 random initial conditions and found that state (c) occurred 515 times and state (d) occurred 485 times. The 95% confidence interval for the probability of state (c) is $p[(c)] \in [0.484, 0.546]$. Here, we cannot conclude that there is any bias toward state (c). But, we can infer that the experimental results are likely skewed toward synchronizing in state (c) compared to the ideal simulation. This could be from a slight experimental asymmetry or a bias in the initial conditions of the phases.

In simulation, we find that for $\alpha = 2$, the results qualitatively remain the same as β increases to at least around $\beta = 1.0$. However, if α is lowered to ~ 0.1, a number of additional states and nonstationary attractors appear in the dynamics for $\beta \gtrsim 0.2$, which could be the subject of further study.

Example 2: Attractive Coupling

We also study multiple ER graphs with reactive-attractive coupling. Fig. 6.12 shows an example of experimental data collected from 200 coupling experiments with the graph shown in Fig. 6.12(a) and the heatmap of all trajectories shown in Fig. 6.12(b). An arbitrary slice of the state space is shown, but all nodes converge to the in-phase synchronized state. In the data shown, $\beta = 0.2$, but we also run experiments with $\beta = 0.6, 1.0$. For this graph, we find the system converges to the in-phase state in all cases for $\beta = 0.2, 0.6$. However, for $\beta = 1.0$, we find that the

system only synchronized to the in-phase state in 193/200 (96.5%) trials. In the remaining 7 trials, the network finds a partially synchronized, nonstationary state.

We collected data for two additional experiments: (I) we constructed a different ER graph using $p = 0.5$; then, with that graph and nodal nonlinearities of $\alpha = 2$, we collected $N_R = 200$ synchronization transitions for 12 values of β ranging from 0.04 to 1.0. We found that the system always goes to the in-phase synchronized state for $\beta \lesssim 0.7$. However, for $\beta \gtrsim 0.8$, we reliably find long lived partially synchronized states of the system. (II) We constructed 20 different ER graphs with $p = 0.5$. We collect $N_R = 100$ synchronization transitions for each graph, with fixed values of $\alpha = 2$ and $\beta = 0.5$. For all of these graphs and all trials, the network synchronizes to the in-phase synchronized state.

More analysis of these experiments will be the subject of future work. We will analyze the existence of these nonstationary states appearing for large values of β. Additionally, we will use the data to analyze the timescales of synchronization in ER random graphs with an eye towards optimizing reservoir computing with nonlinear oscillators.

6.5 Conclusion

We have demonstrated many different synchronized and unsynchronized states which are realized in the same physical network. By varying the nodal operating points and network configuration, we were able to show how different phenomena can be engineered and the states of the network can be controlled. These are basic but fundamental demonstrations of the network's capability, and they are the building blocks from which we aim to achieve nanomechanical computation.

Chapter 7

TOWARDS NEMS NETWORKS WITH ASIC TECHNOLOGY

In order to increase the number of nodes and operating frequency of the network, it is necessary to move beyond PCB technology for interconnecting the devices. Among many possibilities, one promising technique is to integrate the NEMS with an application-specific integrated circuit (ASIC). This can be done via monolithic integration, where the devices are fabricated on the same substrate as the CMOS electronics [120], or heterogeneous integration, where the NEMS are fabricated on a separate chip from the CMOS electronics and interconnected after fabrication. This work focuses on constructing NEMS networks by heterogeneous integration via flip-chip bonding and aims to build a network with 128 oscillators operating at 2 GHz which are coupled in an all-to-all topology. This chapter focuses on the mechanical integration of the two chips and the carrier PCB.

The concept is illustrated by Fig. 7.1, where Fig. 7.1(a) shows the integration plan and Fig. 7.1(b) shows an illustrative schematic of the interconnection strategy. In the CMOS chip, there are 128 independent feedback amplifiers which interface to 128 NEMS devices to create 128 CMOS-NEMS oscillators. These oscillators are interconnected via a global feedback loop within the ASIC which instantiates the all-to-all coupling. Importantly, all of the RF signals are contained inside of the ASIC and NEMS chips, allowing for high-speed operation which is strongly limited by the interconnect parasitic capacitance. However, each device has a third connection, as shown in the figure, which allows for DC or low-frequency biasing of the NEMS devices from an off-chip DAC. Additionally, the ASIC requires a number of voltage rails, current biases, digital inputs, and digital outputs, which are all routed through the NEMS chips and are wirebonded out to the PCB.

This chapter outlines the device fabrication, characterization, and mechanical aspects of the CMOS-NEMS integration.

7.1 GHz NEMS Arrays

The device, pictured in Fig. 7.2(b), is a piezoelectric-ferroelectric lamb-wave resonator. The piezoelectric properties allow the device to efficiently couple electrical inputs and outputs to the mechanical domain, which makes them useful for creating

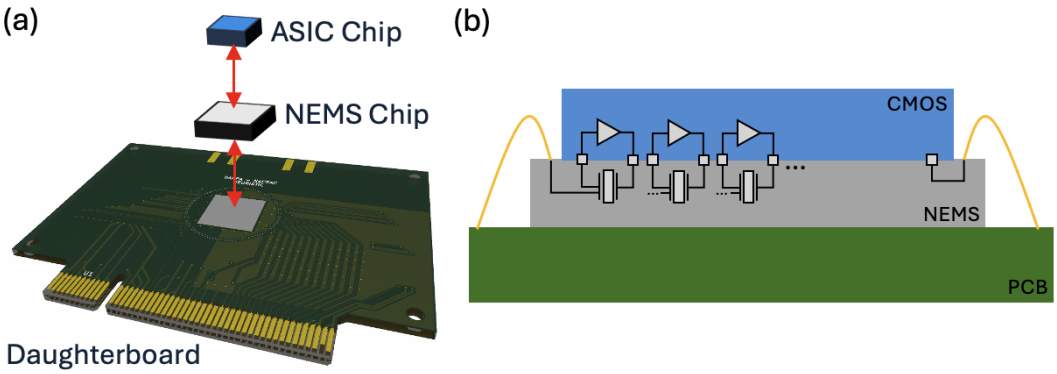

Figure 7.1: ASIC on NEMS Vertical Integration Design. (a) Exploded view of the ASIC-on-NEMS integration approach. (b) Illustrative side view of the chip stack (not to scale). Small squares at the interface between the CMOS chip and NEMS chip represent Au-Au thermosonic bonds. Gold lines on the sides represent wirebonds from the NEMS chip to the PCB. Coupling between nodes is omitted for clarity. All ASIC control signals are passed through the NEMS chip.

high-Q oscillators. The ferroelectric properties allow the device responsivity to be tuned by applying voltage pulses to reset the remanent polarization.

Design and Fabrication

To fabricate the devices, we start with a Si wafer which has a 500 nm layer of 3C-SiC deposited by Nova SiC. Using an Evatec Clusterline 200 II, we sputter 100 nm of AlScN with a 30% Sc concentration. Next, the IDT fingers are defined by a liftoff of a 100 nm Al thin film. The wires, wirebond pads, and device pads are subsequently deposited by liftoff of 200 nm of gold, completing the deposition steps. The release holes are opened by ion beam etching of the AlScN and SiC to the substrate. Finally, the device is suspended by XeF_2 etching of the Si substrate.

We fabricated devices with both gold and aluminum IDT fingers with pitches ranging from 0.4 µm to 10 µm, which yielded devices with fundamental frequencies ranging from approximately 8 GHz to 1 GHz, respectively. We observed a higher yield with the aluminum finger devices than the gold, but the underlying reason is not known. For the integrated system, we required high yield and chose a target frequency of 2 GHz. Thus, based on experiment, we selected the aluminum finger design with a pitch of 6 µm. To match the pads on the ASIC, we put the devices into a 16×8 grid, which is partially shown in Fig. 7.2(a). A zoomed-in view of a single device is shown in Fig. 7.2(b).

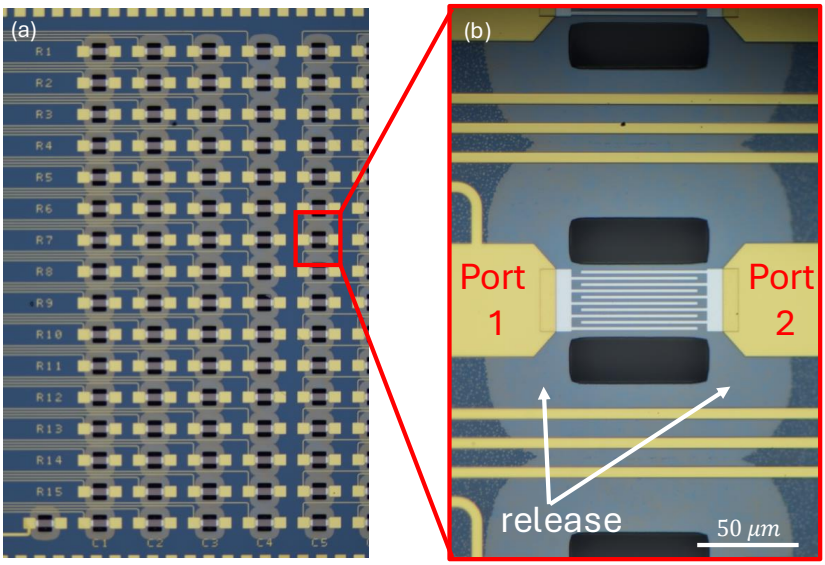

Figure 7.2: NEMS Array. (a) Top-down optical microscope image of the NEMS array. Each device has two bond pads for interfacing with the ASIC as well as an individual wire which supports the DC bias input. (b) Zoomed-in image of a single device. The input and output are labeled as port 1 and port 2, respectively. In the center, the IDT fingers are visible. The released area appears lighter than the unreleased region and is indicated in the figure.

Characterization

To characterize the devices, we use a ground-signal (GS) probe. We land the ground side on port 1 and the signal side on port 2. Then, we perform a reflection measurement to extract the resonance frequency and quality factor. In Fig. 7.3(a) a representative measurement of the reflection magnitude is shown. We see a dip at the mechanical resonance frequency and a peak at the anti-resonance. On this chip, we measured 60 devices and, of those, we achieved 100% yield. Fig. 7.3(b) shows a histogram of the extracted frequencies. The uniformity is very high, with a maximum span of 8.9 MHz. The loaded quality factor for the shown response is $Q_l = 132$, giving a linewidth of 15.0 MHz. Thus, all of the devices fall within a span of ~ 0.59 linewidths.

7.2 Flip-Chip Integration

For this design, we choose the ASIC-on-NEMS approach because the ASIC footprint is very small (2.79 mm by 2.79 mm), so fabricating a NEMS chip to fit on the ASIC is challenging. The small size of the ASIC also sets a limit on the size of the NEMS devices and their bond pads, as the design requires packing 128 oscillators into the system. For this reason, and to minimize parasitic capacitance, the device pads are

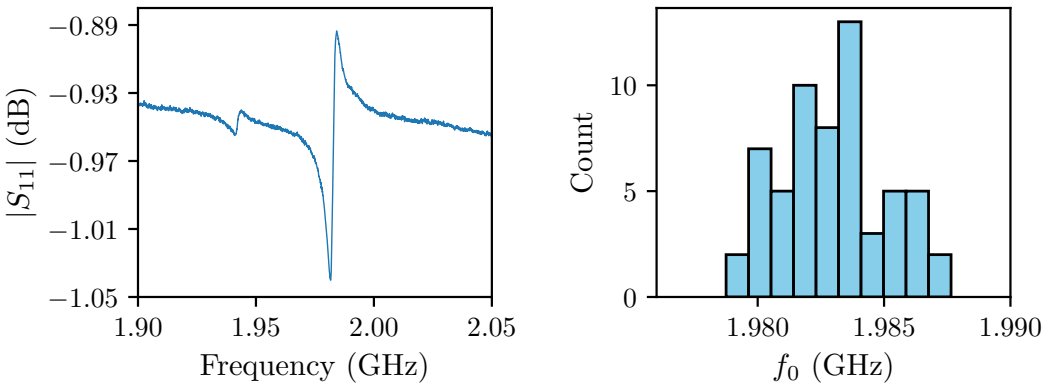

Figure 7.3: NEMS Response and Frequency Histogram. (left) reflection coefficient, $|S_{11}|$, measurement of a typical device. (right) resonant frequency of 60 measured devices.

chosen to be 54 µm by 54 µm and the horizontal and vertical pitch between devices is maximized given the remaining area. Ultimately, this results in an ultra-compact design, fitting 128 devices and their pads into area of approximately 5 mm^2.

The first step of the integration is to attach the NEMS chip to the PCB using silver paint, as shown in Fig. 7.4(a). Then, the ASIC is flip-chip bonded to the NEMS by gold-gold thermosonic bonding. Out of the available flip-chip techniques, this one was chosen because it allows for higher density of interconnects than typical solder bumps; and, anisotropic conductive film (ACF) cannot be placed directly over the suspended NEMS devices. For the bonding, we deposit stud bumps onto the NEMS chip pads using a KS Icon automatic wirebonder. Then, we use a Fineplacer Lambda flip-chip bonder to coin (flatten) the stud bumps before bonding the chips using heat and ultrasonic power. The resulting stack is shown in Fig. 7.4(b) and 7.4(c). For these tests, we fabricated a test version of the ASIC on a transparent substrate (sapphire) so that the alignment and quality could be inspected.

Lastly, 185 wirebond connections are made between the NEMS chip and the PCB, comprising the individual DC bias inputs and the ASIC peripherals. The final assembly is shown in Fig. 7.5.

7.3 Conclusion

The idea of this NEMS-CMOS paradigm is that by combining a NEMS device with a few transistors, it may be possible to build an analog computer which outperforms a purely CMOS or purely mechanical implementation on power, speed, footprint, programmability, or other metrics. This comes at the cost of increased complexity and numerous integration challenges. Here, we demonstrated the device operation

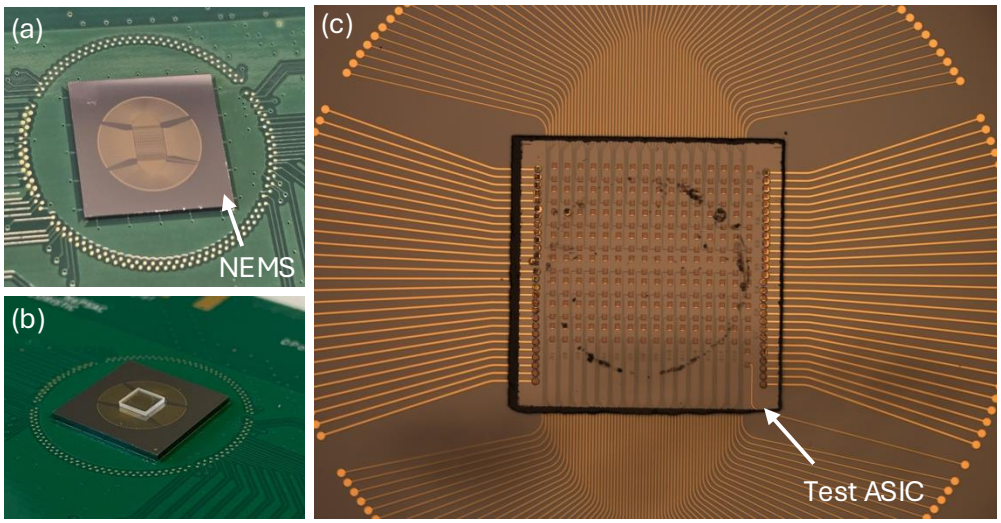

Figure 7.4: ASIC on NEMS Vertical Integration Test. (a) The NEMS chip attached to the PCB, (b) the test ASIC flip-chip bonded to the NEMS chips, and (c) a close up image of the assembled stack. The test ASIC is made of sapphire so the alignment can be inspected.

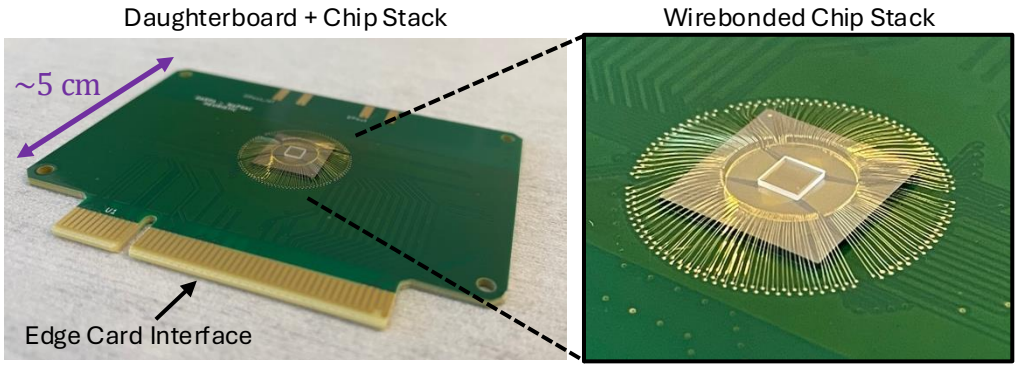

Figure 7.5: Pictures of Mechanical Integration. Pictures of the final test integration stack showing the ASIC-on-NEMS chip stack wirebonded to the PCB.

and mechanical integration. Going forward, more work is needed to address electrical compatibility issues between the ASIC and the NEMS devices to ensure reliable oscillation at 2 GHz.

Chapter 8

CONCLUSION

In this thesis we reviewed the dynamics of nonlinear oscillators and introduced methods to build an analog computer with them. Towards that goal, we developed a wafer-scale fabrication process for highly tunable nonlinear resonators based on thin-film AlN. We integrated those devices into a reconfigurable network of NEMS oscillators with highly tunable nodal nonlinearities and inter-oscillator coupling strengths. By careful study of the dynamical equations of motion, we showed precise calibration and control of the system parameters and demonstrated its capabilities by investigating synchronization and pattern formation on circulant, wheel, and random graphs. Finally, we covered recent progress on designing a GHz frequency, integrated NEMS-ASIC network based on piezoelectric-ferroelectric AlScN resonators for in-memory computing.

Why NEMS?

The physics and methods explored in this thesis are not unique to nanomechanics. Similar dynamics can be engineered in a variety of domains ranging from photonics to superconducting circuits, and similar problems can be solved with other mappings to the physics of those devices. However, as with other complex scientific and engineering endeavors, such as the development of quantum computers, there are a number of complex criteria which must be met and the optimal technological platform is never clear at the outset. In fact, it is usually the case that all of the desirable criteria are not even known.

At present, there is no utility-scale nanomechanical computer; nevertheless, we demonstrated some of the desirable properties of nanomechanics, such as high Q, high tunability, compactness, and strong nonlinearity. We have also exemplified some challenges, such as difficult-to-control material properties, high insertion loss, and slow operating speed. However, we have also argued that these must be considered in light of a specific computing application and that the landscape for computing hardware is fundamentally changing: the CMOS-digital paradigm is reaching its limits and the focus of computing applications is shifting towards AI-driven and edge computing tasks. We envision a future where some computers may actually slow down to better respond to their environment or harness intrinsic

randomness to enhance their capability. For those applications, nanomechanics may be a perfect physical domain for computing.

Future Experiments

In this thesis, we have primarily focused on characterizing steady-state dynamics of the network starting from random initial conditions; this is just the beginning of exploration with this network. Moving forward, we plan to pursue experiments, including (I) exploring exotic states of complex but highly symmetric networks, such as on a cube or directed cyclic graph, (II) exploring dynamics of networks with mixed attractive and repulsive interactions, (III) studies of dynamics under inhomogeneous conditions, (IV) explorations of network control via topological manipulation, (V) realizing explosive synchronization transitions in real-world networks, (VI) reservoir computing with a small network, and (VII) solving 8-node combinatorial optimization problems.

Outlook

On one hand, this project is quite advanced and the bottleneck for moving forward is solving complex engineering challenges. On the other hand, beyond the simplest cases and parameter regimes, it seems apparent that we have hardly scratched the surface of the complex nonlinear dynamics of these high-dimensional systems, particularly as they pertain to real-world experiments. In fact, in the proposed method of reservoir computing, a paradigm based entirely on complexity, we suggested operating the network in the Kuramoto limit, essentially tamping out all of the richness of the higher-order dynamics in exchange for certainty about its operating regime. Or, for the Ising machine, we proposed turning on the coupling infinitely slowly to guarantee finding the global minimum. Are these the ideal methods for computing with these networks? A key challenge going forward will be synthesizing theoretical knowledge of nonlinear networks with practical experimental realizations to move beyond these reductions. We hope these experiments and this platform, along with the many others in the literature, can serve as inspiration and guidance for theorists and experimentalists alike in the pursuit of elucidating the complexity of coupled oscillator networks.

Appendix A

DERIVATION OF SLOW-TIME EQUATIONS OF MOTION FOR WEAKLY COUPLED OSCILLATORS

In this appendix, we derive the slow-time dynamics of a nonlinear oscillator following [29]. We adopt the convention that dimensional variables have a tilde and nondimensional variables have no over-mark. Bar and hat notation is used for intermediate scalings, but all parameters in the final equation are nondimensional. The dimensional equation of motion is given by

$$\tilde{m}\frac{d^2\tilde{x}_n}{d\tilde{t}^2} + \tilde{\gamma}\frac{d\tilde{x}_n}{d\tilde{t}} + \tilde{m}\tilde{\omega}_0^2(1 + \frac{2\tilde{\gamma}}{\tilde{m}\tilde{\omega}_0}\omega_n)\tilde{x}_n + \tilde{\alpha}\tilde{x}_n^3 = \tilde{g}_n G(\tilde{t}) + \tilde{S}_n F(\tilde{x}_n) + \frac{\tilde{\beta}}{2}\sum_{m=1}^{N} C_{nm} H(\tilde{x}_m) \tag{A.1}$$

where x_n[1] is the displacement of node n; $\tilde{t}$ is time; $\tilde{m}$ is the effective mass; $\tilde{\gamma}$ is the linear dissipation; $\tilde{\omega}_0$ is the average natural frequency of all the resonators; ω_n[2] the normalized detuning of the natural frequency of node n; $\tilde{\alpha}$ is the nonlinear stiffness; $\tilde{G}_n$ is an arbitrary weak injection signal with magnitude $\tilde{g}_n$; F is a self-feedback term with strength $\tilde{S}_n$; $\tilde{\beta}$ is a global coupling strength; C_{nm} is the adjacency or graph Laplacian matrix; and $\tilde{H}$ is the coupling function. In practice, G, F, and H can be designed arbitrarily within in the limits of electronically implementable functions, but the specific forms matching the experiment herein will be introduced later.

Next, we scale the equation using two nondimensional variables. Letting $t = \tilde{\omega}_0\tilde{t}$ and $x_n = \tilde{x}_n\sqrt{\tilde{\alpha}_n/\tilde{m}\tilde{\omega}_0^2}$ the equation becomes

$$\ddot{x}_n + \frac{1}{Q}\dot{x}_n + (1 + \frac{2\omega_n}{Q})x_n + x_n^3 = \hat{g}_n G(t) + \hat{S}_n F(x_n, \dot{x}_n) + \frac{\hat{\beta}}{2}\sum_{m=1}^{N} C_{nm} H(x_m, \dot{x}_m) \tag{A.2}$$

where dot notation refers to derivative with respect to nondimensional time, t, and the nondimensional parameters are given by

[1] Throughout this work, we assume linear transduction between the electrical and mechanical domain, meaning that this equation can be written for $\tilde{x}$ in units of displacement or voltage as they have a linear relationship determined by the electromechanical responsivity.

[2] We chose ω_n to denote the nondimensional frequency detuning so that the final equation's notation is generally consistent with the dynamical systems literature.

$$Q^{-1} = \frac{\tilde{\gamma}}{\tilde{m}\omega_0};\;\; \hat{g}_n = \tilde{g}_n \frac{1}{\omega_0^3}\sqrt{\frac{\tilde{\alpha}_n}{\tilde{m}^3}};\;\; \hat{S}_n = \tilde{S}_n \frac{1}{\omega_0^3}\sqrt{\frac{\tilde{\alpha}_n}{\tilde{m}^3}};\;\; \hat{\beta} = \tilde{\beta}\frac{1}{\omega_0^3}\sqrt{\frac{\tilde{\alpha}_n}{\tilde{m}^3}}. \tag{A.3}$$

Note on Frequency Detuning Term

Similar to [27], in (A.1) frequency inhomogeneity is introduced through the node dependent parameter ω_n. The form is chosen for two reasons: (1) it will appear as a second-order correction in the perturbation solution and (2) ω_n will approximately equate to the frequency detuning of resonator n from the average, normalized to the 3 dB linewidth of the device. Here we explore why. From (A.1), the linear natural frequency of resonator n is given by

$$\tilde{\omega}_n^2 = \tilde{\omega}_0^2\left(1 + \frac{2\tilde{\gamma}}{\tilde{m}\tilde{\omega}_0}\omega_n\right) = \tilde{\omega}_0^2\left(1 + \frac{2\omega_n}{Q}\right). \tag{A.4}$$

Solving for ω_n yields

$$\omega_n = \frac{Q}{2}\left(\frac{\tilde{\omega}_n^2}{\tilde{\omega}_0^2} - 1\right) = \frac{Q}{2}\left(\frac{\tilde{f}_n^2}{\tilde{f}_0^2} - 1\right) \tag{A.5}$$

Now, let $\tilde{f}_n = \tilde{f}_0 + \Delta\tilde{f}$ where $\Delta\tilde{f} \ll \tilde{f}_0$ is the small detuning. Therefore, $\tilde{f}_n^2 = \tilde{f}_0^2 + 2\tilde{f}_0\Delta\tilde{f} + \Delta\tilde{f}^2$. Neglecting the $\Delta\tilde{f}^2$ term and plugging this back into (A.5) yields

$$\omega_n \approx Q\frac{\tilde{f}_0\Delta\tilde{f}}{\tilde{f}_0^2} = Q\frac{\tilde{f}_n - \tilde{f}_0}{\tilde{f}_0} = \frac{\tilde{f}_n - \tilde{f}_0}{\Delta f_{3\text{dB}}} \tag{A.6}$$

which shows that under this construction, ω_n is approximately equal to the dimensional frequency detuning divided by the linewidth. This approximation is a good assumption for small detuning. For example, with $f_0 = 1.8$ MHz, $Q = 2500$, the linewidth is $\Delta f_{3\text{dB}} = 720$ Hz. At a detuning of 10 linewidths (7.2 kHz), the exact formula gives $\omega_n = 10.02$ and the approximation gives $\omega_n = 10.00$, a relative error of $\approx 0.2\%$. Table A.1 shows the absolute and relative error for a few select values of detuning. Depending on the application, the error may need to be accounted for. But, for applications in this work, we deal with small detuning and define ω_n by (A.6).

Table A.1: Exact vs. Approximate ω_n at Selected Detunings and with Illustrative Parameter Values ($f_0 = 1.8\,\text{MHz}$, $Q = 2500$, $\Delta f_{3\text{dB}} = 720\,\text{Hz}$).

Detuning	$\Delta\tilde{f}$ (Hz)	$\omega_{n,\text{exact}}$	$\omega_{n,\text{approx}}$	Abs. err.	Rel. err. (%)
$1\,\Delta f_{3\text{dB}}$	720.0	1.0002	1.0000	0.0002	0.0200
$5\,\Delta f_{3\text{dB}}$	3600.0	5.0050	5.0000	0.0050	0.0999
$10\,\Delta f_{3\text{dB}}$	7200.0	10.0200	10.0000	0.0200	0.1996
$20\,\Delta f_{3\text{dB}}$	14400.0	20.0800	20.0000	0.0800	0.3984

A.1 Solution via Secular Perturbation Theory

To solve this coupled nonlinear equation, we use secular perturbation theory [29], commonly referred to as the method of multiple scales or two-timing. The devices have very low dissipation, so we define a small parameter, $\epsilon \equiv 1/Q << 1$ which is used to construct a perturbative solution. But, the solution should still be constrained by our initial assumption that the restoring force can be truncated at the third order in displacement. Therefore, we impose the requirement that the device response should be such that $x \propto \sqrt{\epsilon}$. This ensures that $x^3 \propto \epsilon^{3/2}$, a factor of ϵ smaller. From this, we know that the feedback strength should be on the order of $\epsilon^{3/2}$. The other physical effects, such as external injection and coupling can be on this order or smaller. With these conditions, we expand $x_n(t)$ and its derivatives as

$$x_n(t) = \epsilon^{1/2} x_n^{(0)}(t) + \epsilon^{3/2} x_n^{(1)}(t) + \epsilon^{5/2} x_n^{(2)}(t) + \ldots \tag{A.7}$$

$$\dot{x}_n(t) = \epsilon^{1/2} \dot{x}_n^{(0)}(t) + \epsilon^{3/2} \dot{x}_n^{(1)}(t) + \epsilon^{5/2} \dot{x}_n^{(2)}(t) + \ldots \tag{A.8}$$

$$\ddot{x}_n(t) = \epsilon^{1/2} \ddot{x}_n^{(0)}(t) + \epsilon^{3/2} \ddot{x}_n^{(1)}(t) + \epsilon^{5/2} \ddot{x}_n^{(2)}(t) + \ldots \tag{A.9}$$

The superscripts in parentheses indicate each order of the expansion while the subscript, n, still indicates the oscillator number. In the following subsection, we define each term explicitly.

Coupling

We consider linear coupling, which can be generically written as a combination of reactive and dissipative coupling as

$$\frac{\hat{\beta}}{2} \sum_{m=1}^{N} C_{nm} H(x_m, \dot{x}_m) = \frac{\hat{\beta}}{2} \sum_{m=1}^{N} C_{nm} \left[\lambda x_m + \zeta \dot{x}_m\right] \tag{A.10}$$

where $\lambda \in \mathbb{R}$ and $\zeta \in \mathbb{R}$ are the $O(1)$ reactive and dissipative coefficients, respectively. In the model (and in the experiment), C_{nm} can be defined as an adjacency matrix, whereby the diagonal terms are $C_{nn} = 0\ \forall n$. Or, as the graph Laplacian, whereby the diagonals are equal to the degree of the node and off-diagonals are given be the negative of the adjacency matrix. Alternatively, graph Laplacian coupling is achieved by writing the coupling as

$$\frac{\hat{\beta}}{2}\sum_{m=1}^{N} C_{nm}\left[\lambda(x_m - x_n) + \zeta(\dot{x}_m - \dot{x}_n)\right] \tag{A.11}$$

where C_{nm} is defined as the adjacency matrix. For more details see [4]. The form of the coupling can affect the dynamics. Hereafter, we use the form of (A.10) and do not assume a form of the coupling matrix.

Later, we will see that by allowing β to be complex, we can capture reactive and dissipative coupling by changing the phase. Thus, we drop the dissipative part of the coupling for the remainder of the derivation and absorb λ into $\hat{\beta}$.

Imposing the magnitude condition of the expansion, we write the final coupling term as

$$\epsilon^{3/2}\frac{\bar{\beta}}{2}\sum_{m=1}^{N} C_{nm}x_m \tag{A.12}$$

where we have explicitly pulled out $\epsilon^{3/2}$ such that $\bar{\beta} = \hat{\beta}\epsilon^{-3/2}$ and $\bar{\beta} \sim O(1)$.

Direct Injection Signal

For direct injection, the signal should be close to the oscillation frequency to induce effects such as frequency pulling. Here, we assume that the injection for oscillator n is a sinusoidal signal with amplitude, g_n; scaled frequency, $1 + \epsilon\Omega_n$; and a tunable phase shift, Δ_n. Ω_n describes small deviations of the injected frequency from the intrinsic oscillation frequency, which is 1 in the scaled time. The injection signal written in complex notation has the form of (A.13).

$$\hat{g}_n G(t) = \epsilon^{3/2}\bar{g}_n e^{i(\epsilon\Omega_n t + \Delta_n)}e^{it} \tag{A.13}$$

where $\epsilon^{3/2}$ is explicitly pulled out such that $\bar{g}_n$ is $O(1)$.

Feedback

The feedback is a function which is dependent on both the amplitude and frequency of the oscillation and so cannot be written down compactly in the second-order displacement equations. The feedback contains three major components: (1) saturated feedback is achieved with an AGC amplifier, which uses an internal power detector and control loop to maintain a constant output. (2) A bandpass filter reduces harmonics generated by the nonlinearity. And, (3) a phase shifter adjusts the feedback to be in-phase with velocity. In the displacement domain, we can write down a generalized function capturing these effects

$$F(x_n, t) = \frac{\mathcal{F}[x_n](t)}{\mathcal{P}[x_n](t)} = \frac{\int_{-\infty}^{t} h(t-\tau)\, x_n(\tau)\, d\tau}{\frac{1}{\sqrt{2}}\sqrt{\int_{-\infty}^{t} \mathcal{K}(t-\tau)\, x_n^2(\tau)\, d\tau}} \tag{A.14}$$

where $\mathcal{F}$ is a linear time-invariant (LTI) operator which applies a phase shift and bandpass filter to the signal, x_n, described by the impulse response, h. In the denominator, $\mathcal{P}$ is a nonlinear time-invariant operator which computes the signal amplitude based on a kernel, $\mathcal{K}$, which is designed to emulate the effect of the power detector and gain control loop in the AGC circuit. This form is not readily amenable to analytical treatment. However, by assuming the AGC response time is much shorter than the timescale of the envelope variations of the oscillator, we can treat the amplitude as quasi-static within the AGC amplifier's integration window. Thus, the operators act on the signal as if the envelope was constant, eliminating the explicit time dependence of the feedback. (A.14) becomes

$$F(x_n) = \frac{\int_{-\infty}^{0} h(-\tau)\, x_n(\tau)\, d\tau}{\frac{1}{\sqrt{2}}\sqrt{\int_{-\infty}^{0} \mathcal{K}(-\tau)\, x_n^2(\tau)\, d\tau}}. \tag{A.15}$$

In practice, the AGC circuit has a nominal step response settling time of 30 µs, which is > 10× faster than the amplitude relaxation time constant of $\tau_a = Q/\pi f_0 = 442$ µs. Under these assumptions and by treating the signal as a single tone, we can avoid treating this analytically and instead introduce a saturated feedback of the form in (A.25) to the slow-time envelope equations of motion.

Solution

Putting everything together, we have an equation describing the network:

$$\ddot{x}_n + \epsilon\dot{x}_n + (1 + 2\epsilon\omega_n)x_n + x_n^3 =$$
$$\epsilon^{3/2}S_nF(x_n, \dot{x}_n) + \epsilon^{3/2}\bar{g}_n e^{i(\epsilon\Omega_n t+\Delta_n)}e^{it} + \epsilon^{3/2}\frac{\bar{\beta}}{2}\sum_{m=1}^{N} C_{nm}x_m. \quad \text{(A.16)}$$

Plugging the perturbation expansion (A.7), (A.8), and (A.9), the lowest order solution ($O(\epsilon^{1/2})$) comes out to be the simple harmonic oscillator equation:

$$\ddot{x}_n^{(0)}(t) + x_n^{(0)}(t) = 0 \quad \text{(A.17)}$$

which has the solution

$$x_n^{(0)}(t) = \frac{1}{2}e^{it} + \frac{1}{2}e^{-it}. \quad \text{(A.18)}$$

The problem arises when we go for the next order in the expansion. In the regular approach to perturbation theory, we would put (A.18) into the expansion, (A.7), and solve for the next order, which is hopefully enough to capture the dynamics for small enough ϵ. We get a solution of the form

$$\ddot{x}_n^{(1)} + x_n^{(1)} = [...]e^{it} + [...]e^{i3t}. \quad \text{(A.19)}$$

However, this describes a simple harmonic oscillator with resonant forcing term, e^{it} (highlighted in red). Thus, the solution will contain secular terms of the form te^{it} which diverge as $t \to \infty$. This cannot be a physical solution, so something went wrong. See Section 7.6 of [1] for a full discussion, but the basic problem is that there are two time scales in the system and the perturbation is to the frequency of the device, not its displacement. Regular perturbation fails in this case; in order to get a solution accurate for long times, we would need to take an infinite number of terms in the expansion. Alternatively, we employ secular perturbation theory, which is also known as two-timing or the method of multiple scales. Here, we will allow the lowest order solution to have a slowly varying amplitude and phase described by a complex amplitude, $A(T)$, where $T = \epsilon t$ is the slow time scale of the system. This complex amplitude varies much slower than the oscillation frequency of the system — so much slower that we can treat A as constant on the time scale of the oscillations.

Plugging in the lowest order solution and allowing the amplitude and phase to vary slowly yields the updated perturbation expansion

$$x_n(t) = \frac{\epsilon^{1/2}}{2}\left[\tilde{A}_n(T)e^{it} + c.c.\right] + \epsilon^{3/2}x_n^{(1)}(t) + \epsilon^{5/2}x_n^{(2)}(t) + ... \tag{A.20}$$

with $\tilde{A}_n(T) = \tilde{a}_n e^{i\phi_n}$. $\tilde{a}_n$ and ϕ_n are the amplitude and phase of the n^{th} oscillator, respectively. We use a tilde in these equations because $\tilde{A}$ will be scaled. These values describe the envelope of the oscillating signal and can be extracted from experimental data using the Hilbert transformation. We plug this expansion into (A.16) using

$$\dot{\tilde{A}} = \frac{d\tilde{A}}{dt} = \frac{d\tilde{A}}{dT}\frac{dT}{dt} = \epsilon\frac{d\tilde{A}}{dT} = \epsilon\tilde{A}' \tag{A.21}$$

and

$$\ddot{\tilde{A}} = \epsilon^2\tilde{A}'' \tag{A.22}$$

which yields, via product rule, the derivatives

$$\dot{x}_n(t) = \frac{\epsilon^{1/2}}{2}\left[(\epsilon\tilde{A}_n' + i\tilde{A}_n)e^{it} + c.c.\right] + \epsilon^{3/2}\dot{x}_n^{(1)}(t) + ... \tag{A.23}$$

and

$$\ddot{x}_n(t) = \frac{\epsilon^{1/2}}{2}\left[(\epsilon^2\tilde{A}_n'' + i2\epsilon\tilde{A}_n' - \tilde{A}_n)e^{it} + c.c.\right] + \epsilon^{3/2}\ddot{x}_n^{(1)}(t) + ... \tag{A.24}$$

Additionally, following [119] and as discussed in the feedback section, we introduce the following expression for fully saturated feedback:

$$\tilde{S}_n F(\tilde{x}_n) = \epsilon^{3/2}\frac{S_n}{2}e^{i\tilde{\phi}_n}e^{i\xi}e^{it} = \epsilon^{3/2}\frac{S_n}{2}\frac{\tilde{A}_n}{|\tilde{A}_n|}e^{i\xi}e^{it} \tag{A.25}$$

where S_n is the saturation amplitude for node n, ϕ_n is the phase of oscillator n, ξ is the phase shift of the feedback, and e^{it} represents the signal coming from the resonator to lowest order. This term represents the a feedback signal which is pinned at a specific amplitude (independent of $\tilde{a}_n$) and which has certain phase shift, ξ. If

we plug (A.20), (A.23), and (A.24) into (A.16), apply the feedback form (A.25), and group terms of order $\epsilon^{3/2}$, we get

$$\ddot{x}_n^{(1)} + x_n^{(1)} = \left[-i\tilde{A}_n' - i\frac{1}{2}\tilde{A}_n - \omega_n\tilde{A}_n - \frac{3}{8}|\tilde{A}_n|^2\tilde{A}_n\right] e^{it} + \bar{g}_n e^{i\epsilon\Omega_n t + i\Delta_n} e^{it} + \frac{S_n}{2}\frac{\tilde{A}_n}{|\tilde{A}_n|}e^{i\xi}e^{it} + \bar{\beta}\sum_{m=1}^{N} C_{nm}\tilde{A}_m e^{it} + [\ldots]e^{i3t}.$$

The structure of this solution has the same problem — the highlighted red terms are all resonant forcing terms which will lead to secular solutions. To avoid these terms, we require that the coefficients of the secular forcing terms must sum to zero for all time. This yields a differential equation for the slow time complex amplitude:

$$\frac{d\tilde{A}_n}{dT} = -i\frac{S_n}{2}\frac{\tilde{A}_n}{|\tilde{A}_n|}e^{i\xi} - \frac{1}{2}\tilde{A}_n + i\omega_n\tilde{A}_n + i\frac{3}{8}|\tilde{A}_n|^2\tilde{A}_n - i\bar{g}_n e^{i\Delta_n}e^{i\epsilon\Omega_n t} - i\frac{\bar{\beta}}{2}\sum_{m=1}^{N} C_{nm}\tilde{A}_m.$$

We set the phase of the feedback to be purely dissipative such that $\xi = \pi/2$. We scale the complex amplitude by S_n such that $\tilde{A}_n = S_n A_n$, and we choose to arbitrarily shift the coupling phase by $e^{i\pi/2}$, giving

$$\frac{dA_n}{dT} = \frac{1}{2}\frac{A_n}{|A_n|} - \frac{1}{2}A_n + i\left(\omega_n + \alpha_n|A_n|^2\right)A_n - ig_n e^{i\Delta_n}e^{i\epsilon\Omega_n t} - i\frac{\beta}{2}\sum_{m=1}^{N} C_{nm}A_m$$

where we defined the parameter $\alpha_n = \frac{3}{8}S_n^2$ and performed the final scaling of $g_n = \bar{g}_n/S_n$ and $\beta = \bar{\beta}e^{i\pi/2}/S_n$. In this thesis, we typically consider the network without external drive signals, so we set $g_n = 0$ and consider the autonomous dynamics given by

$$\frac{dA_n}{dT} = \frac{1}{2}\frac{A_n}{|A_n|} - \frac{1}{2}A_n + i\left(\omega_n + \alpha_n|A_n|^2\right)A_n + \frac{1}{2}\beta\sum_{m=1}^{N} C_{nm}A_m \qquad \text{(A.26)}$$

where a complex value of β allows the equation to capture dissipative and reactive coupling. Above, we specifically chose the phase of the coupling term so that as

it is written here, $\beta = \zeta - i\lambda$ defines dynamics where ζ is the dissipative coupling strength and λ is the reactive coupling strength. And, positive ζ and λ correspond to attractive coupling for each type. We follow this convention throughout this work.

Appendix B

DERIVATION OF HIGHER ORDER PHASE DYNAMICS IN THE SMALL COUPLING LIMIT

B.1 Intro

In experiment, many configurations of the coupling can be realized. Here we derive the higher-order dynamics for purely dissipative and purely reactive coupling of uniform strength but for any network connectivity. Starting from (A.26) and following the same conventions, the equations of motion for a network of coupled oscillators with quadratic dispersive nonlinearity, saturated feedback, and linear coupling are given by

$$\frac{dA_n}{dT} = \frac{1}{2}\frac{A_n}{|A_n|} - \frac{1}{2}A_n + i\left(\omega_n + \alpha_n|A_n|^2\right)A_n + \frac{1}{2}(\zeta - i\lambda)\sum_{m=1}^{N} C_{nm}A_m \tag{B.1}$$

where ζ is the magnitude of the dissipative coupling and λ is the magnitude of the reactive coupling. As it is written, both coefficients are attractive when positive and repulsive when negative. $C_{nm} \in \{0, 1\}$ is the adjacency matrix of the network.

To examine the amplitude and phase dynamics of the system, we split the equations of motion into real and imaginary parts giving a coupled set of first-order differential equations

$$a_n' = \frac{1}{2} - \frac{1}{2}a_n + \frac{\zeta}{2}\sum_{m=1}^{N} C_{nm}a_m\cos(\phi_m - \phi_n) + \frac{\lambda}{2}\sum_{m=1}^{N} C_{nm}a_m\sin(\phi_m - \phi_n) \tag{B.2}$$

$$\phi_n' = \omega_n + \alpha_n a_n^2 + \frac{\zeta}{2}\sum_{m=1}^{N} C_{nm}\frac{a_m}{a_n}\sin(\phi_m - \phi_n) - \frac{\lambda}{2}\sum_{m=1}^{N} C_{nm}\frac{a_m}{a_n}\cos(\phi_m - \phi_n). \tag{B.3}$$

In the next sections, we limit the dynamics to reactive and dissipative coupling and explore the difference in phase dynamics.

B.2 Dissipative Coupling

For dissipative coupling, we set $\lambda = 0$, giving

$$a_n' = \frac{1}{2} - \frac{1}{2}a_n + \frac{\zeta}{2}\sum_{m=1}^{N} C_{nm}a_m \cos(\phi_m - \phi_n) \tag{B.4}$$

$$\phi_n' = \omega_n + \alpha_n a_n^2 + \frac{\zeta}{2}\sum_{m=1}^{N} C_{nm}\frac{a_m}{a_n}\sin(\phi_m - \phi_n). \tag{B.5}$$

Asymptotic Expansion

Here, we carry out the asymptotic expansion of the amplitude-phase dynamics for the dissipative coupling case. For dissipative coupling, the amplitude-phase equations are given by B.4 and B.5. We expand the amplitude and phase in powers of ζ as follows

$$a_n = a_n^{(0)} + \zeta a_n^{(1)} + \zeta^2 a_n^{(2)} + O(\zeta^3) \tag{B.6}$$

$$\phi_n = \phi_n^{(0)} + \zeta \phi_n^{(1)} + \zeta^2 \phi_n^{(2)} + O(\zeta^3) \tag{B.7}$$

thus

$$a_n' = a_n^{(0)'} + \zeta a_n^{(1)'} + \zeta^2 a_n^{(2)'} + O(\zeta^3) \tag{B.8}$$

$$\phi_n' = \phi_n^{(0)'} + \zeta \phi_n^{(1)'} + \zeta^2 \phi_n^{(2)'} + O(\zeta^3). \tag{B.9}$$

The procedure from here is to plug Eqs. B.6-B.9 into Eqs. B.4 and B.5 and solve the equations at each order in ζ. Starting with the amplitude, we get the following equation

$$\begin{aligned} a_n^{(0)'} + \zeta a_n^{(1)'} + \zeta^2 a_n^{(2)'} + O(\zeta^3) = {} & \frac{1}{2} - \frac{1}{2}a_n^{(0)} \\ & + \zeta\left[-\frac{1}{2}a_n^{(1)} + \frac{1}{2}\sum_{m}^{N} C_{nm}a_m^{(0)}\cos(\phi_m - \phi_n)\right] \\ & + \zeta^2\left[-\frac{1}{2}a_n^{(2)} + \frac{1}{2}\sum_{m}^{N} C_{nm}a_m^{(1)}\cos(\phi_m - \phi_n)\right] + O(\zeta^3) \end{aligned}$$

at zeroth order

$$a_n^{(0)'} = \frac{1}{2} - \frac{1}{2}a_n^{(0)}. \tag{B.10}$$

We want to solve for steady-state solutions in amplitude, so we set $a_n^{(0)'} = 0$ and find

$$a_n^{(0)} = 1. \tag{B.11}$$

Collecting terms of order ζ, we have

$$a_n^{(1)'} = -\frac{1}{2}a_n^{(1)} + \frac{1}{2}\sum_m^N C_{nm}a_m^{(0)}\cos(\phi_m - \phi_n) \tag{B.12}$$

again, setting $a_n^{(1)'} = 0$ and plugging in the solution for $a_m^{(0)}$ yields

$$a_n^{(1)} = \sum_m^N C_{nm}\cos(\phi_m - \phi_n) \tag{B.13}$$

and, finally, at order ζ^2 we find

$$a_n^{(2)'} = -\frac{1}{2}a_n^{(2)} + \frac{1}{2}\sum_m^N C_{nm}a_m^{(1)}\cos(\phi_m - \phi_n). \tag{B.14}$$

setting $a_n^{(2)'} = 0$ and inserting the solution for $a_m^{(1)}$ gives

$$a_n^{(2)} = \sum_{m,k}^N C_{nm}C_{mk}\cos(\phi_k - \phi_m)\cos(\phi_m - \phi_n) \tag{B.15}$$

Now, we have the solution for the amplitude up to $O(\zeta^2)$. So, we turn our attention to the phase expansion. Plugging in (B.7) into (B.5) gives

$$\begin{aligned}\phi_n^{(0)'} + \zeta\phi_n^{(1)'} + \zeta^2\phi_n^{(2)'} &= \omega_n + \alpha_n\left[a_n^{(0)} + \zeta a_n^{(1)} + \zeta^2 a_n^{(2)} + O\left(\zeta^3\right)\right]^2 \\ &\quad + \frac{\zeta}{2}\sum_m^N C_{nm}\left[\frac{a_m^{(0)} + \zeta a_m^{(1)} + \zeta^2 a_m^{(2)} + O\left(\zeta^3\right)}{a_n^{(0)} + \zeta a_n^{(1)} + \zeta^2 a_n^{(2)} + O\left(\zeta^3\right)}\sin(\phi_m - \phi_n)\right].\end{aligned}$$

To proceed, we use the Taylor series expansion

$$\frac{1}{1+x} = 1 + x + x^2 + O\left(x^3\right) \tag{B.16}$$

for small values of x. Noting that the coupling term has a ζ multiplying the entire term, we only need to expand up to $O(\zeta)$ here

$$\begin{aligned}
\frac{1}{a_n^{(0)} + \zeta a_n^{(1)} + O\left(\zeta^2\right)} &= \frac{1}{a_n^{(0)}} \frac{1}{1 + \zeta \frac{a_n^{(1)}}{a_n^{(0)}} + O\left(\zeta^2\right)} \\
&= \frac{1}{a_n^{(0)}} + \frac{1}{a_n^{(0)}} \left[\zeta \frac{a_n^{(1)}}{a_n^{(0)}} + O\left(\zeta^2\right) \right] + O\left(\zeta^2\right) \\
&= \frac{1}{a_n^{(0)}} + \zeta \frac{a_n^{(1)}}{a_n^{(0)^2}} + O\left(\zeta^2\right)
\end{aligned}$$

which, after some algebra, gives the following equation for ϕ_n:

$$\begin{aligned}
&\phi_n^{(0)'} + \zeta\, \phi_n^{(1)'} + \zeta^2 \phi_n^{(2)'} + O\left(\zeta^3\right) \\
&\quad = \omega_n + \alpha_n a_n^{(0)2} \\
&\qquad + \zeta \left[2\alpha_n a_n^{(0)} a_n^{(1)} + \frac{1}{2} \sum_m^N C_{nm} \frac{a_m^{(0)}}{a_n^{(0)}} \sin(\phi_m - \phi_n) \right] \\
&\qquad + \zeta^2 \Bigg[\alpha_n a_n^{(1)2} + 2\alpha_n a_n^{(0)} a_n^{(2)} \\
&\qquad\qquad + \frac{1}{2} \sum_m^N C_{nm} \left(-\frac{a_m^{(0)} a_n^{(1)}}{a_n^{(0)2}} \sin(\phi_m - \phi_n) + \frac{a_m^{(1)}}{a_n^{(0)}} \sin(\phi_m - \phi_n) \right) \Bigg] \\
&\qquad + O\left(\zeta^3\right)
\end{aligned}$$

plugging in our known solutions for $a_n^{(0)}$, $a_n^{(1)}$, and $a_n^{(2)}$, we can solve for each order of the phase equations

$$\phi_n^{(0)'} = \omega_n + \alpha_n \tag{B.17}$$

$$\phi_n^{(1)'} = 2\alpha_n \sum_m^N C_{nm} \cos(\phi_m - \phi_n) + \frac{1}{2} \sum_m^N C_{nm} \sin(\phi_m - \phi_n) \tag{B.18}$$

$$\begin{aligned}
\phi_n^{(2)'} &= \alpha_n \left[\sum_m^N C_{nm} \cos(\phi_m - \phi_n) \right]^2 \\
&+ 2\alpha_n \sum_{m,k}^N C_{nm} C_{mk} \cos(\phi_k - \phi_m) \cos(\phi_m - \phi_n) \\
&- \frac{1}{2} \sum_{m,k}^N C_{nm} C_{nk} \cos(\phi_k - \phi_n) \sin(\phi_m - \phi_n) \\
&+ \frac{1}{2} \sum_{m,k}^N C_{nm} C_{mk} \cos(\phi_k - \phi_m) \sin(\phi_m - \phi_n).
\end{aligned} \tag{B.19}$$

After some substantial algebra, the second-order term can be rewritten as

$$\begin{aligned}
\phi_n^{(2)'} &= \frac{1}{2} \sum_m C_{nm}^2 \left[\alpha_n + \alpha_n \cos\left[2(\phi_m - \phi_n)\right] - \frac{1}{2} \sin[2(\phi_m - \phi_n)] \right] \\
&+ \frac{1}{2} \sum_m C_{nm} C_{mm} \left[4\alpha_n \cos(\phi_m - \phi_n) + \sin(\phi_m - \phi_n) \right] \\
&+ \frac{1}{4} \sum_m C_{nm} \sum_{k \neq m} C_{nk} \Big\{ 2\alpha_n \cos(\phi_k - \phi_m) + 2\alpha_n \cos(\phi_k + \phi_m - 2\phi_n) \\
&\qquad - \sin(\phi_k + \phi_m - 2\phi_n) + \sin(\phi_k - \phi_m) \Big\} \\
&+ \frac{1}{4} \sum_m C_{nm} \sum_{k \neq m} C_{mk} \Big\{ 4\alpha_n \cos(\phi_k - \phi_n) + 4\alpha_n \cos(\phi_k + \phi_n - 2\phi_m) \\
&\qquad - \sin(\phi_k + \phi_n - 2\phi_m) + \sin(\phi_k - \phi_n) \Big\}.
\end{aligned}$$

Combining all the terms, the second-order phase dynamics for dissipative coupling are given by (B.20):

$$
\begin{aligned}
\phi_n' &= \omega_n + \alpha_n \\
&+ \frac{\zeta}{2}\sum_m^N C_{nm}\Big\{\sin(\phi_m-\phi_n) + 4\alpha_n\cos(\phi_m-\phi_n)\Big\} \\
&+ \frac{\zeta^2}{2}\sum_m C_{nm}C_{mm}\left[4\alpha_n\cos(\phi_m-\phi_n) + \sin(\phi_m-\phi_n)\right] \\
&+ \frac{\zeta^2}{4}\sum_m C_{nm}^2\left[2\alpha_n + 2\alpha_n\cos\left[2(\phi_m-\phi_n)\right] - \sin[2(\phi_m-\phi_n)]\right] \\
&+ \frac{\zeta^2}{4}\sum_m C_{nm}\sum_{k\neq m} C_{nk}\Big\{2\alpha_n\cos(\phi_k-\phi_m) + 2\alpha_n\cos(\phi_k+\phi_m-2\phi_n) \\
&\qquad\qquad - \sin(\phi_k+\phi_m-2\phi_n) + \sin(\phi_k-\phi_m)\Big\} \\
&+ \frac{\zeta^2}{4}\sum_m C_{nm}\sum_{k\neq m} C_{mk}\Big\{4\alpha_n\cos(\phi_k-\phi_n) + 4\alpha_n\cos(\phi_k+\phi_n-2\phi_m) \\
&\qquad\qquad - \sin(\phi_k+\phi_n-2\phi_m) + \sin(\phi_k-\phi_n)\Big\}.
\end{aligned}
\tag{B.20}
$$

B.3 Reactive Coupling

For reactive coupling, we set $\zeta = 0$, giving the amplitude and phase equations

$$a_n' = \frac{1}{2} - \frac{1}{2}a_n + \frac{\lambda}{2}\sum_{m=1}^N C_{nm}a_m\sin(\phi_m-\phi_n) \tag{B.21}$$

$$\phi_n' = \omega_n + \alpha_n a_n^2 - \frac{\lambda}{2}\sum_{m=1}^N C_{nm}\frac{a_m}{a_n}\cos(\phi_m-\phi_n). \tag{B.22}$$

From here, we can write down an asymptotic expansion for the dynamics in powers of λ.

$$a_n = a_n^{(0)} + \lambda a_n^{(1)} + \lambda^2 a_n^{(2)} + O\left(\lambda^3\right) \tag{B.23}$$

$$\phi_n = \phi_n^{(0)} + \lambda\phi_n^{(1)} + \lambda^2\phi_n^{(2)} + O\left(\lambda^3\right). \tag{B.24}$$

Next, we plug equations Eqs. B.23 and B.24 into Eqs. B.21 and B.22. The steps follow exactly those shown in Section B.2, so we only show the key results here.

First, the amplitude terms are given by

$$a_n^{(0)} = 1 \tag{B.25}$$

$$a_n^{(1)} = \sum_m^N C_{nm} \sin(\phi_m - \phi_n) \tag{B.26}$$

$$a_n^{(2)} = \sum_{m,k}^N C_{nm} C_{mk} \sin(\phi_k - \phi_m) \sin(\phi_m - \phi_n). \tag{B.27}$$

Second, the phase terms are given by

$$\phi_n^{(0)'} = \omega_n + \alpha_n \tag{B.28}$$

$$\phi_n^{(1)'} = 2\alpha_n \sum_m^N C_{nm} \sin(\phi_m - \phi_n) - \frac{1}{2} \sum_m^N C_{nm} \cos(\phi_m - \phi_n) \tag{B.29}$$

$$\begin{aligned} \phi_n^{(2)'} = {} & \alpha_n \left\{ \sum_m C_{nm} \sin(\phi_m - \phi_n) \right\}^2 \\ & + 2\alpha_n \sum_{m,k} C_{nm} C_{mk} \sin(\phi_k - \phi_m) \sin(\phi_m - \phi_n) \\ & + \frac{1}{2} \sum_{m,k} C_{nm} C_{nk} \sin(\phi_k - \phi_n) \cos(\phi_m - \phi_n) \\ & - \frac{1}{2} \sum_{m,k} C_{nm} C_{mk} \sin(\phi_k - \phi_m) \cos(\phi_m - \phi_n). \end{aligned} \tag{B.30}$$

Thus, the phase dynamics up to second-order in λ are

$$
\begin{aligned}
\phi_n' &= \omega_n + \alpha_n \\
&+\frac{\lambda}{2}\sum_m^N C_{nm}\Big\{4\alpha_n \sin(\phi_m - \phi_n) - \cos(\phi_m - \phi_n)\Big\} \\
&+\frac{\lambda^2}{4}\sum_m C_{nm}^2\Big\{2\alpha_n - 2\alpha_n \cos\left[2(\phi_m - \phi_n)\right] + \sin[2(\phi_m - \phi_n)]\Big\} \\
&+\frac{\lambda^2}{4}\sum_m\sum_{k\neq m} C_{nm}C_{nk}\Big\{\, 2\alpha_n \cos(\phi_k - \phi_m) - 2\alpha_n \cos(\phi_k + \phi_m - 2\phi_n) \\
&\qquad\qquad + \sin(\phi_k + \phi_m - 2\phi_n) + \sin(\phi_k - \phi_m)\Big\} \\
&+\frac{\lambda^2}{4}\sum_m\sum_{k\neq m} C_{nm}C_{mk}\Big\{ - 4\alpha_n \cos(\phi_k - \phi_n) + 4\alpha_n \cos(\phi_k + \phi_n - 2\phi_m) \\
&\qquad\qquad - \sin(\phi_k + \phi_n - 2\phi_m) - \sin(\phi_k - \phi_n)\Big\}.
\end{aligned}
\tag{B.31}
$$

Appendix C

DYNAMICS OF TWO COUPLED OSCILLATORS

Here, we examine the equations of motion for dissipative and reactive coupling with $N = 2$. Here, we limit the discussion to first-order phase dynamics.

C.1 Dissipative Coupling

Starting from (B.20), we let $N = 2$ and

$$\mathbf{C} = \begin{bmatrix} 0 & C_{12} \\ C_{21} & 0 \end{bmatrix}$$

which yields a set of four equations

$$a_1 = 1 + \zeta C_{12} \cos(\phi_2 - \phi_1) \tag{C.1}$$

$$a_2 = 1 + \zeta C_{21} \cos(\phi_1 - \phi_2) \tag{C.2}$$

$$\phi_1' = \omega_1 + \alpha_1 + 2\alpha_1 \zeta C_{12} \cos(\phi_2 - \phi_1) + \frac{1}{2}\zeta C_{12} \sin(\phi_2 - \phi_1) \tag{C.3}$$

$$\phi_2' = \omega_2 + \alpha_2 + 2\alpha_2 \zeta C_{21} \cos(\phi_1 - \phi_2) + \frac{1}{2}\zeta C_{21} \sin(\phi_1 - \phi_2). \tag{C.4}$$

We can further reduce the equations by solving for the phase difference. Applying the change of variables $\phi = \phi_2 - \phi_1$ and $\phi' = \phi_2' - \phi_1'$ we get

$$a_1 = 1 + \zeta C_{12} \cos(\phi) \tag{C.5}$$

$$a_2 = 1 + \zeta C_{21} \cos(\phi) \tag{C.6}$$

$$\phi' = \Delta\omega + \Delta\alpha + 2\zeta(C_{21}\alpha_2 - C_{12}\alpha_1) \cos(\phi) - \frac{1}{2}\zeta(C_{21} + C_{12}) \sin(\phi) \tag{C.7}$$

where $\Delta\omega = \omega_2 - \omega_1$ and $\Delta\alpha = \alpha_2 - \alpha_1$. If $\alpha_1 = \alpha_2 = \alpha$, ϕ' reduces to

$$\phi' = \Delta\omega + 2\alpha\beta(C_{21} - C_{12}) \cos(\phi) - \frac{1}{2}\zeta(C_{21} + C_{12}) \sin(\phi). \tag{C.8}$$

Furthermore, if the coupling is symmetric such that $C_{12} = C_{21} = 1$, the equations reduce to their final form

$$a_1 = 1 + \zeta \cos(\phi) \tag{C.9}$$
$$a_2 = 1 + \zeta \cos(\phi) \tag{C.10}$$
$$\phi' = \Delta\omega - \zeta \sin(\phi). \tag{C.11}$$

This set of equations describes two dissipatively coupled oscillators which (1) are in the small coupling limit ($\zeta \ll 1$), (2) have identical amplitude-frequency coupling, α, and (3) have symmetric interaction strength. It is clear that such oscillators synchronize purely as a function of the coupling strength, ζ and do not depend on any nonlinearity. There are two fixed points in the dynamics, regardless of the sign of ζ. But, for $\zeta > 0$, only the in-phase synchronized state is stable whereas for $\zeta < 0$, the out-of-phase synchronized state is stable. Thus, as derived here, $\zeta > 0$ is attractive (ferromagnetic) while $\zeta < 0$ is repulsive (antiferromagnetic).

We can analytically compute the synchronization conditions as a function of ζ and $\Delta\omega$ by plugging in the phase locking condition, $\phi' = 0$, into (C.11).

$$\Delta\omega = \zeta \sin(\phi) \tag{C.12}$$

which only has a solution if

$$|\Delta\omega| \leq |\zeta|. \tag{C.13}$$

Thus, (C.13) gives the range of $\Delta\omega$ and ζ under which the oscillators frequency-lock. The specific phase difference in the synchronization regime is given by

$$\phi = \sin^{-1}\left(\frac{\Delta\omega}{\zeta}\right). \tag{C.14}$$

Plugging this solution into (C.9) and (C.10) produces the predicted amplitudes in the synchronized state

$$a_1 = a_2 = 1 - \zeta \cos\left[\sin^{-1}\left(\frac{\Delta\omega}{\zeta}\right)\right] \tag{C.15}$$

recalling that $\cos(\arcsin(x)) = \sqrt{1 - x^2}$, the amplitudes can be rewritten as

$$a_1 = a_2 = 1 - \zeta\sqrt{1 - \left(\frac{\Delta\omega}{\zeta}\right)^2}. \tag{C.16}$$

Outside of the synchronization regime, the amplitude equation does not hold. But, for the phase dynamics described by (C.11) we can solve the expression of the average frequency difference, $\langle\phi'\rangle$.

$$\left\langle\frac{d\phi}{dt}\right\rangle = \sqrt{\Delta\omega^2 - \zeta^2} \tag{C.17}$$

Collecting these results, Fig. C.1 show the expected amplitude and frequency as a function of $\Delta\omega$ for multiple values of ζ.

C.2 Reactive Coupling

Starting from (B.31) and letting

$$\mathbf{C} = \begin{bmatrix} 0 & C_{12} \\ C_{21} & 0 \end{bmatrix}$$

yields

$$a_1 = 1 + \lambda C_{12}\sin(\phi_2 - \phi_1) \tag{C.18}$$

$$a_2 = 1 + \lambda C_{21}\sin(\phi_1 - \phi_2) \tag{C.19}$$

$$\phi_1' = \omega_1 + \alpha_1 + 2\alpha_1\lambda C_{12}\sin(\phi_2 - \phi_1) - \frac{1}{2}\lambda C_{12}\cos(\phi_2 - \phi_1) \tag{C.20}$$

$$\phi_2' = \omega_2 + \alpha_2 + 2\alpha_2\lambda C_{21}\sin(\phi_1 - \phi_2) + \frac{1}{2}\lambda C_{21}\cos(\phi_1 - \phi_2) \tag{C.21}$$

letting $C_{12} = C_{21} = 1$; $\alpha_1 = \alpha_2 = \alpha$; $\phi = \phi_1 - \phi_2$; and $\Delta\omega = \omega_1 - \omega_2$:

$$a_1 = 1 - \lambda\sin(\phi) \tag{C.22}$$

$$a_2 = 1 + \lambda\sin(\phi) \tag{C.23}$$

$$\phi' = \Delta\omega - 4\alpha\lambda\sin(\phi). \tag{C.24}$$

The synchronization condition is given by

$$|\Delta\omega| \leq |4\alpha\lambda| \tag{C.25}$$

and the phase difference in the synchronization regime is

$$\phi = \sin^{-1}\left(\frac{\Delta\omega}{4\alpha\lambda}\right) \tag{C.26}$$

and the amplitudes are given by

$$a_{1,2} = 1 \mp \frac{\Delta\omega}{4\alpha} \tag{C.27}$$

and, finally, the average frequency outside the synchronization regime is given by

$$\left\langle \frac{d\phi}{dt} \right\rangle = \sqrt{\Delta\omega^2 - 16\alpha^2\lambda^2}. \tag{C.28}$$

C.3 Reactive vs. Dissipative Coupling

For two mutually coupled oscillators in the small coupling limit, the dynamics are summarized in Table C.1 and plotted for select parameter values in Fig. C.1. We can see that the form of phase dynamics is identical for both cases, that is, both reactive and dissipative coupling yield a sine coupling between the phases. Thus, we can get dissipative phase dynamics from both dissipatively coupled and reactively coupled oscillators. However, reactively coupled oscillators have a 4α coefficient on the coupling, which can significantly boost or inhibit the coupling strength for strong or weak nonlinearity. In experiment, α is highly tunable with an ultimate device limited range from nearly $\alpha = 0$ to $\alpha > 100$.

Table C.1: Summary of Dynamics of Two Coupled Oscillators

Dynamics	Dissipative	Reactive
Phase	$\phi' = \Delta\omega - \zeta \sin(\phi)$	$\phi' = \Delta\omega - 4\alpha\lambda \sin(\phi)$
Amplitude	$a_{1,2} = 1 + \zeta \cos(\phi)$	$a_{1,2} = 1 \mp \lambda \sin(\phi)$

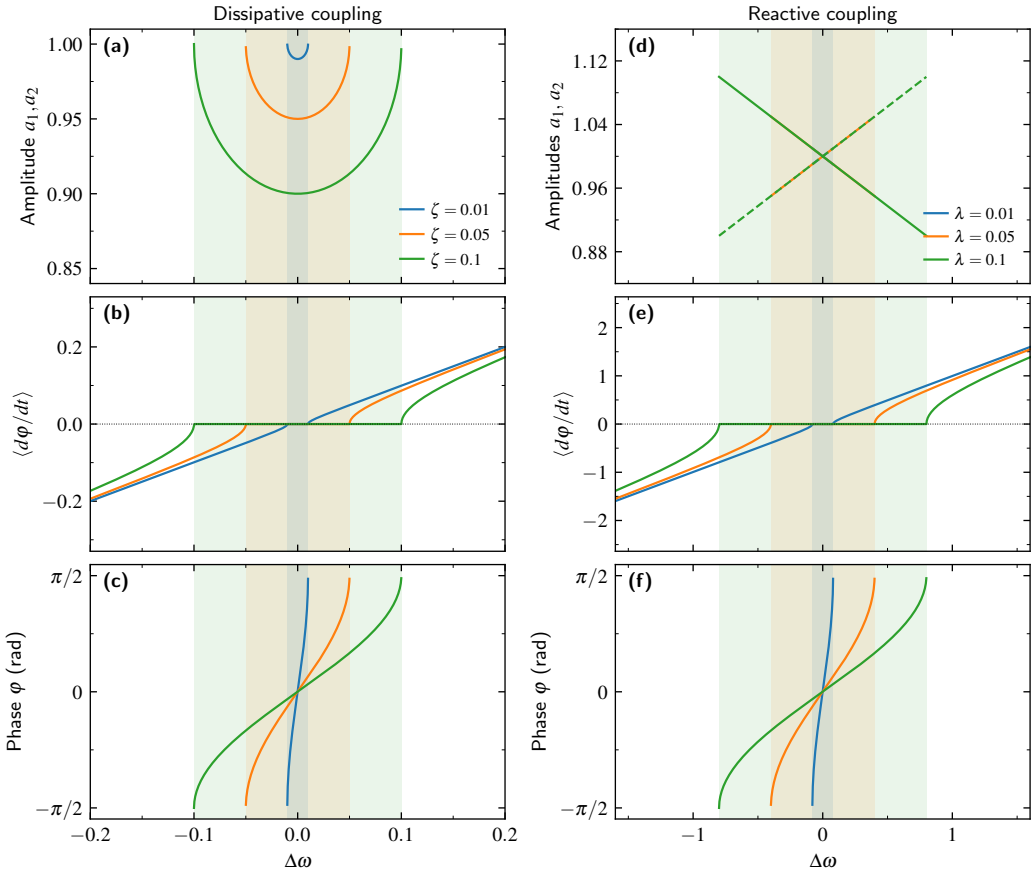

Figure C.1: Locking Dynamics of Two Coupled Oscillators. For $\alpha = 2$, (top) steady-state amplitude in the synchronization regime, (middle) average frequency outside of the synchronization regime, (bottom) phase difference inside the synchronization regime. All functions are plotted for (left) dissipative coupling and (right) reactive coupling.

Appendix D

ADDITIONAL DEVICE CHARACTERIZATION

In the main body of the text, we presented measurements characterizing the primary electromechanical properties of the resonator relevant to the main experiments. Here, we present supplementary characterization data, such as amplitude sweeps with wider ranges, raw frequency response data as a function of DC bias, and additional chip and wafer level statistics.

D.1 Large Amplitude Nonlinear Response

Using the measurement setup in Fig. 4.7, the device response is measured with drive powers from -80 dBm (0.01 nW) up to -6 dBm (0.25 mW). The response amplitude is always the value measured at the lock-in amplifier, so it includes the gain from the readout amplifier.

In Fig. D.1 the amplitude and phase response are plotted starting from -54 dBm drive power. Lower drive powers are omitted so that the phase response is not obscured by noisy traces. We see the same characteristics as described in §4.3. In this case, we see that the device is stable up to very large nonlinear frequency shifts. At -6 dBm, the frequency is shifted by ∼85 kHz (∼132 linewidths).

To further exemplify the character of the response, the same response is plotted with a logarithmic scale on the y-axis in Fig. D.2. Additionally, this plot includes data for drive powers as low as -80 dBm. If the drive power is reduced significantly more, the resonance response is not observed. For the lowest drive power, the noise floor is observed for drive powers away from resonance. With this measurement setup (probe station) and these measurement parameters (bandwidth, averaging, etc...), the thermomechanical noise could be measured.

In Fig. D.3, for the same resonant response and data, the forward transmission coefficient is plotted. S21 is calculated by $S21 = 20\log_{10}(V_{meas}/V_{drv}) - G$, where $G = 40$ is the gain of the readout amplifier in dB. This is an approximation of the S21. Calibrated measurements should be done to extract the precise values; however, the effect of nonlinear damping can still be observed by examining the relative change in S21 on resonance. There is a clear drop in transmission on resonance as the drive power is increased. Characterizing the nonlinear damping is beyond the scope

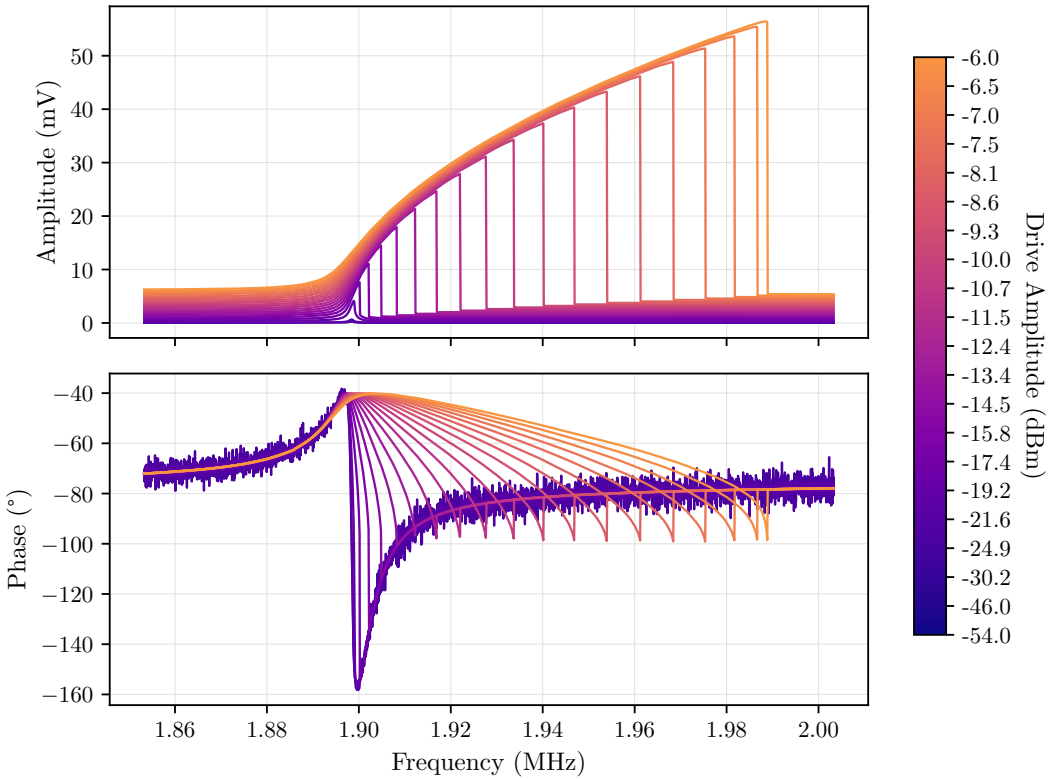

Figure D.1: Resonator Response at Large Drive Power (Linear Amplitude). Amplitude and phase response of a typical resonator measured from -54 dBm to -6 dBm drive power. Frequency is swept upward.

of this work, but experiments demonstrate that it can be ignored in the regime of operation of the oscillator network.

D.2 Frequency Response vs DC Bias

In Fig. D.4, frequency and amplitude sweeps are plotted for -5 V, 0 V, and +5 V. For clarity, we only plot these three bias points. From this raw data, we extract the backbone curves and nonlinear coefficients as a function of the bias. The backbone curves for all of the bias points are plotted in Fig. D.5. Here, we plot the data as frequency shift normalized by the linewidth at that particular DC bias value to eliminate the effect of the changing linear frequency and Q. We extract the backbone curves as a function of output voltage because it accounts for changes in electromechanical coupling over the DC bias range. Using the output voltage gives a true sense of the change in the duffing coefficient because it is proportional to displacement whereas the drive amplitude is conflated with electromechanical coupling and Q changes. We observed a small decrease in the nonlinear coefficient with increasing DC bias.

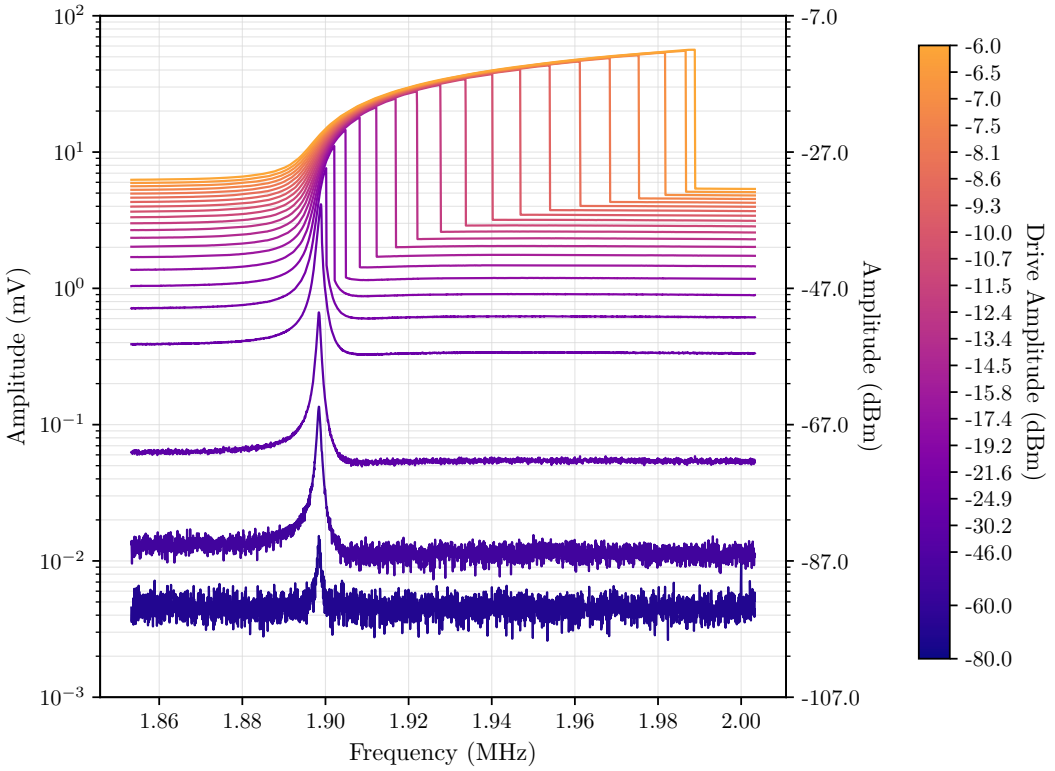

Figure D.2: Resonator Response at Large Drive Power (Logarithmic). Amplitude response of the resonator plotted on a logarithmic amplitude scale. The response amplitude is shown in mV and dBm on the left and right y-axes, respectively. This plot includes the response for all of the measured drive powers which are as low as -80 dBm (0.01 nW). Frequency is swept upward.

D.3 Die and Wafer Level Statistics

In Fig. D.6 we report the device statistics by device number. On a particular chip, there are 15 devices which linearly vary in size from ~40×40 µm to ~45×45 µm. As designed, we observed a consistent frequency downshift as the devices get larger. Additionally, Fig. D.7 shows histograms of the linear frequency and Q for all devices measured. The distributions do not indicate individual device sizes or chips but do give a general sense of the overall variations.

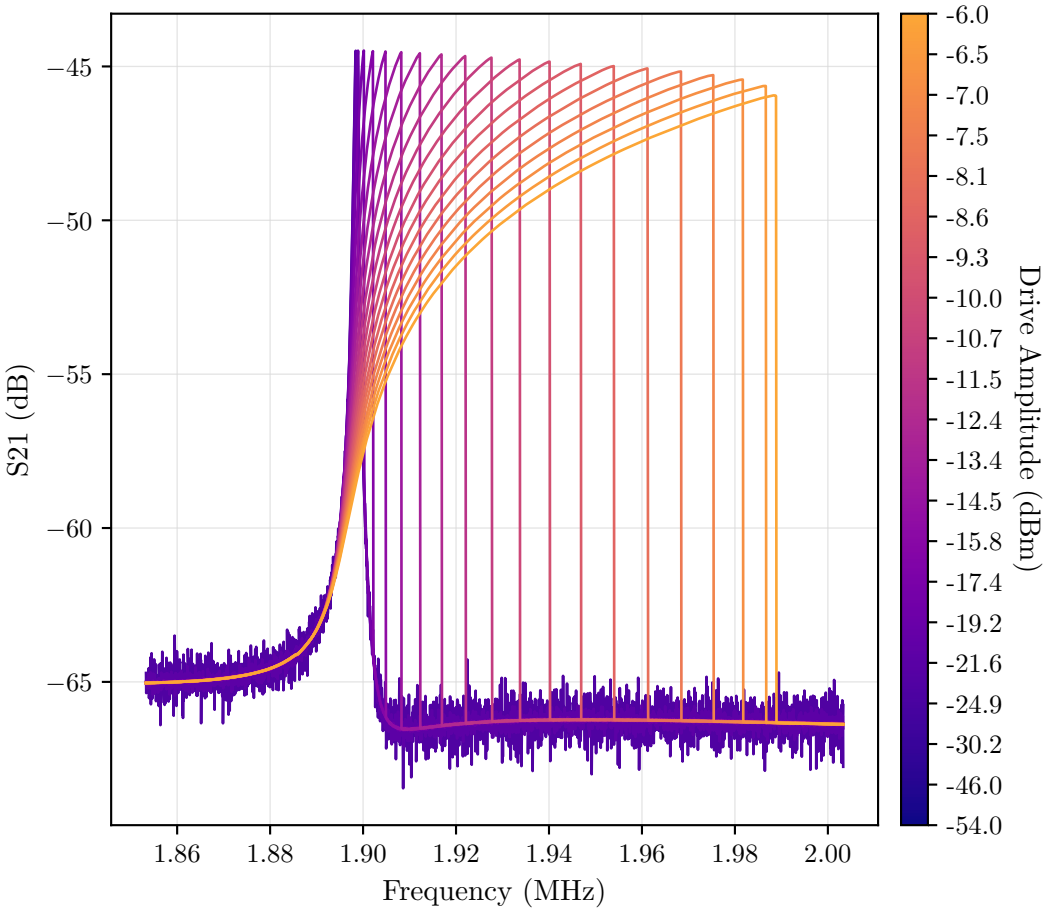

Figure D.3: Resonator Transmission (S21). Plot of the forward transmission, S21, of the device as the power and frequency are swept. Frequency is swept upward.

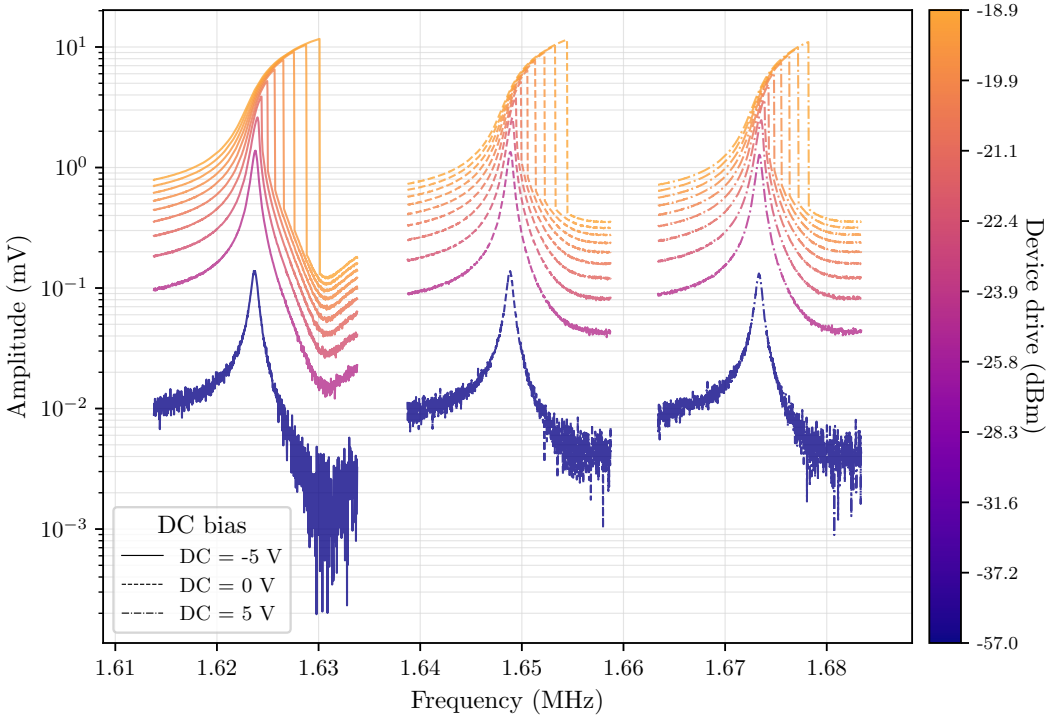

Figure D.4: Raw Data of DC and Amplitude Sweep. Frequency and amplitude sweeps for a resonator at -5 V, 0 V, and +5 V DC bias points. In all cases, frequency is swept upward.

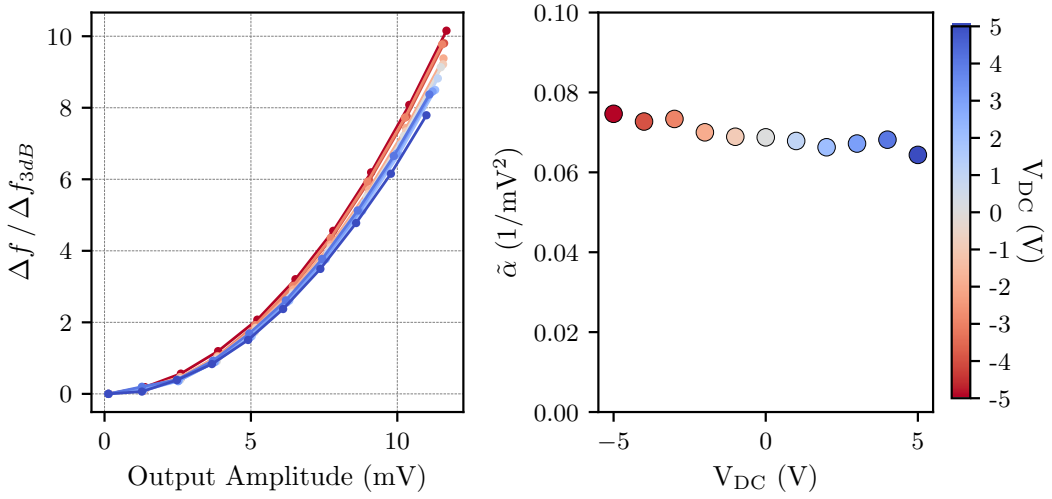

Figure D.5: Measured Backbone Curves for each DC Bias. (a) Measured backbone curves for bias voltage from -5 V to 5 V in 1 V increments. (b) Extracted nonlinear stiffness coefficient for each bias level. $\tilde{\alpha}$ is plotted in units of linewidths per mV^2.

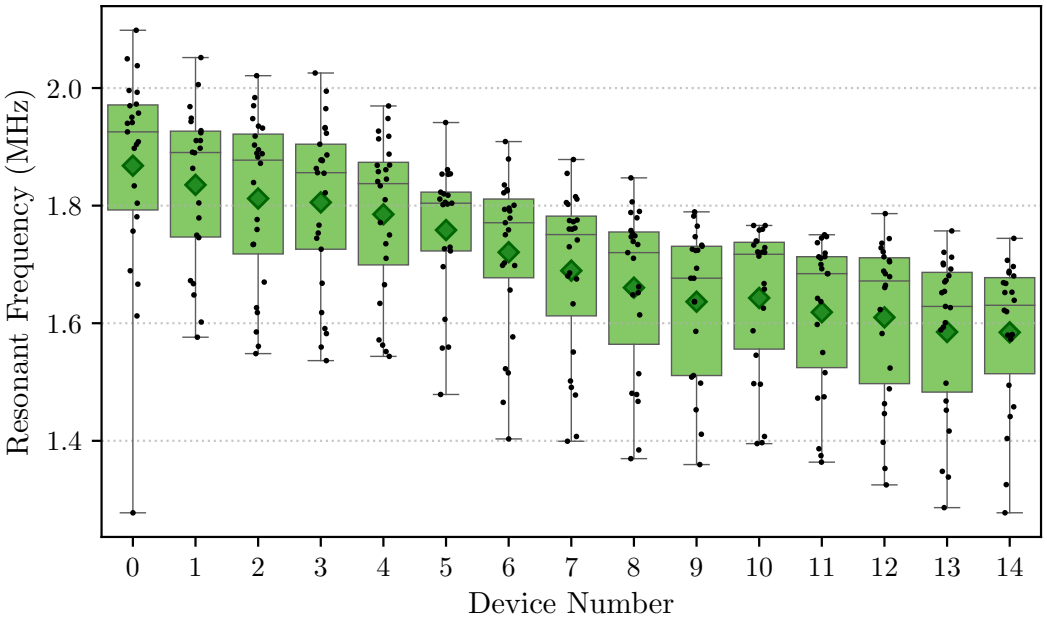

Figure D.6: Resonant Frequency vs Device Size. Statistics for all measured devices as a function of the device size. From device 0 to device 14, the size linearly decreases from ~40×40 µm to ~45×45 µm. Black dots show the frequency of individual devices. Dark green diamonds indicate the average frequency for that device number. Green histograms show the median, upper quartile, and lower quartile.

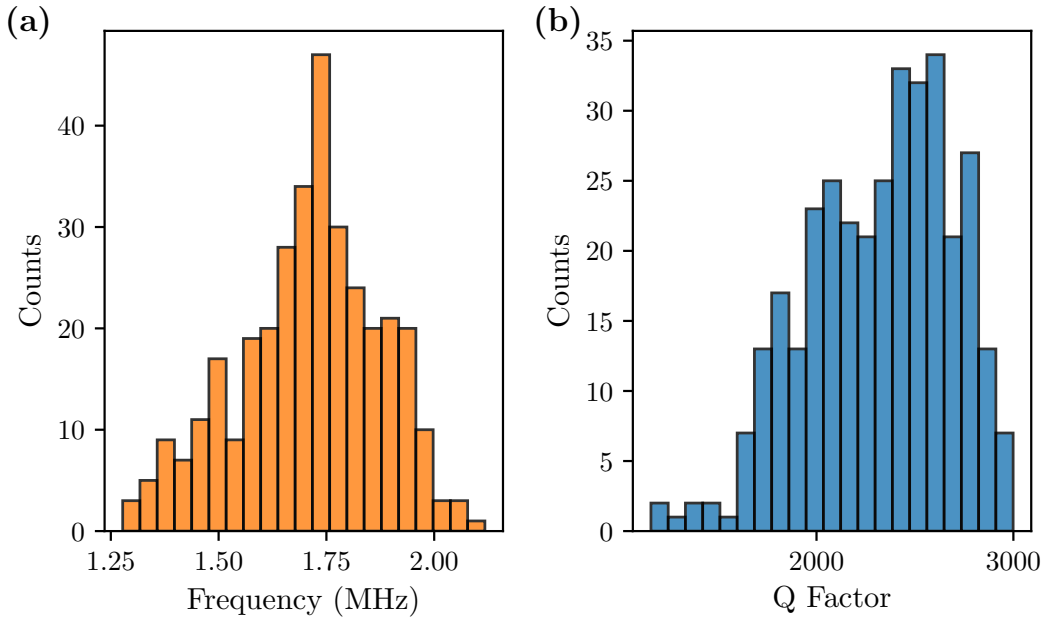

Figure D.7: Frequency and Q Histograms. Histograms of (a) linear resonant frequency and (b) Q for all devices measured.

Appendix E

DETAILED EXPERIMENTAL SETUP

E.1 Hardware Overview

Front and back isometric photos of the experimental setup are shown in Fig. E.1. The network is constructed using standard, commercially available PCB technology. This allows for fast and inexpensive prototyping compared to integration with ASIC technology, and it is feasible for networks with 2 to $\sim$ 100 nodes. The network described here is modular. The nodes and edges are part of separate PCBs which are interconnected by a motherboard.

Firstly, the nodes of the network, labeled (B) in Fig. E.1, are credit-card sized PCBs which each contain a single NEMS resonator embedded into a feedback loop. These PCBs are 6 layers and measure 10x6 cm. A close up image of the oscillator PCB is shown in Fig. E.2. The circuit board is fabricated with an electroless-nickel electroless-palladium immersion-gold (ENEPIG) surface finish so that the NEMS can be directly wirebonded to the circuit. This is advantageous over other techniques, such as using a chip carrier, because it minimizes the parasitic capacitance and footprint. The NEMS is encapsulated by a local vacuum chamber which is 3D printed for this application. Each vacuum chamber is connected to a manifold which is pumped using a Pfeiffer vacuum pumping station. The vacuum level reached is <100 mTorr, which is sufficiently low to effectively eliminate the air damping and saturate the Q. Each oscillator contains a number of amplifiers, attenuators, and other ICs to enable precise control of the operating point of the oscillator over a wide range.

Secondly, the coupling between the nodes is controlled by the“network card”, labeled (A) in Fig. E.1. Through a 98 pin PCIe edge card, the network card connects to the motherboard. It enables control over the topology of the network and the interactions strengths, as described in Chapter 5.

Finally, the motherboard PCB is on the bottom. It houses and integrates all of the other components through PCIe edge connections. Its primary function is to transfer RF signals between the nodes via the network card, but it also serves to route input and output signals, supply DC power, and distribute digital control lines.

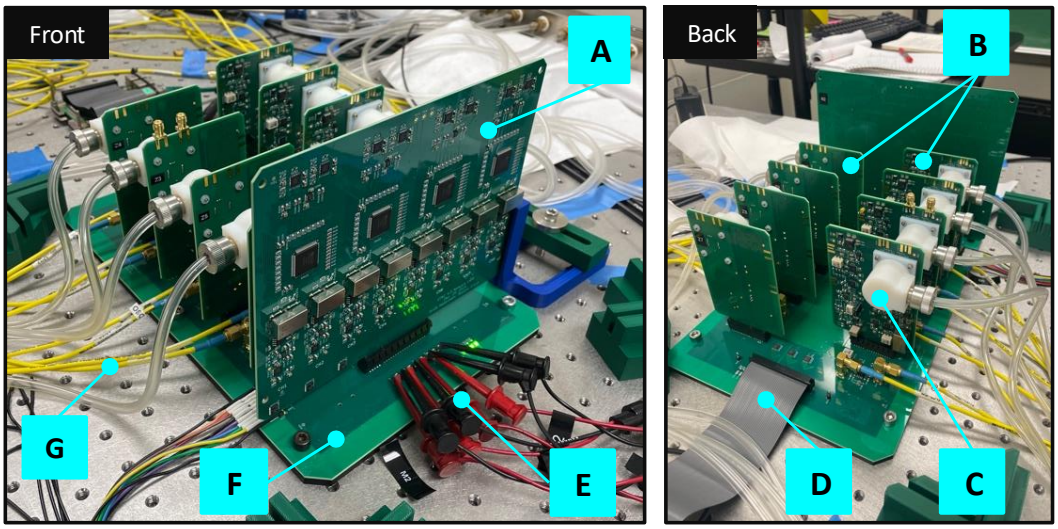

Figure E.1: Experiment Hardware. Detailed experimental setup: isometric views of the front and back of the experiment. (A) Edges: modular, 8-channel programmable network card. (B) Nodes: modular oscillator cards. Each one contains one device wirebonded into a controllable feedback loop. (C) nodal vacuum chambers encapsulating each NEMS chip. (D) Ribbon cable connection to the network for digital communication. (E) DC power supply for all network components. (F) Motherboard connecting all the modular components and distributing RF signals, digital signals, and power. (G) SMA cables providing RF input (injection signals) and RF output (nodal measurement) to each node in the network.

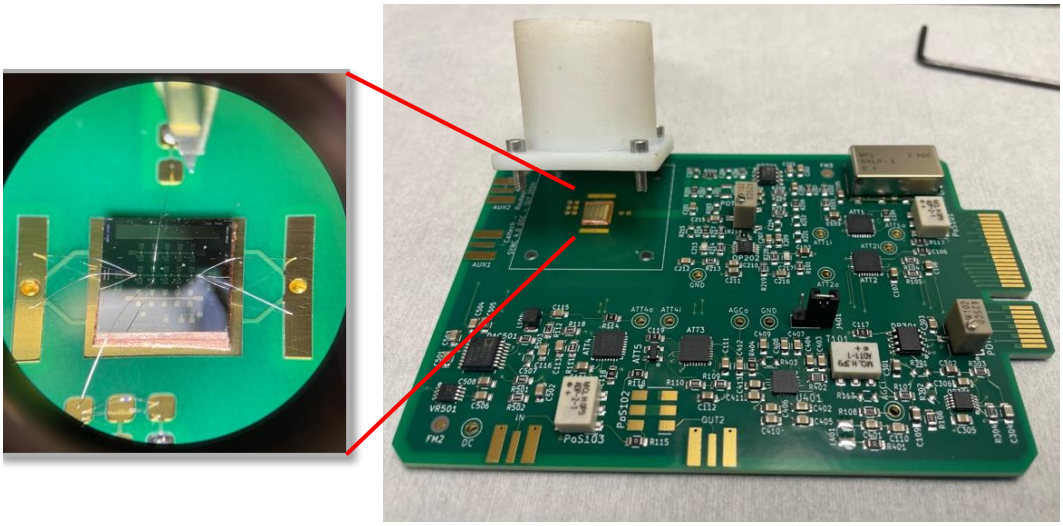

Figure E.2: Oscillator PCB. (right) picture of the oscillator PCB. (left) image of the NEMS die from which one device is bonded into the oscillator circuit. The electronic feedback loop includes numerous components for controlling the oscillator's setpoint (effective nonlinearity, amplitude, frequency) and for analog interfacing to the other nodes.

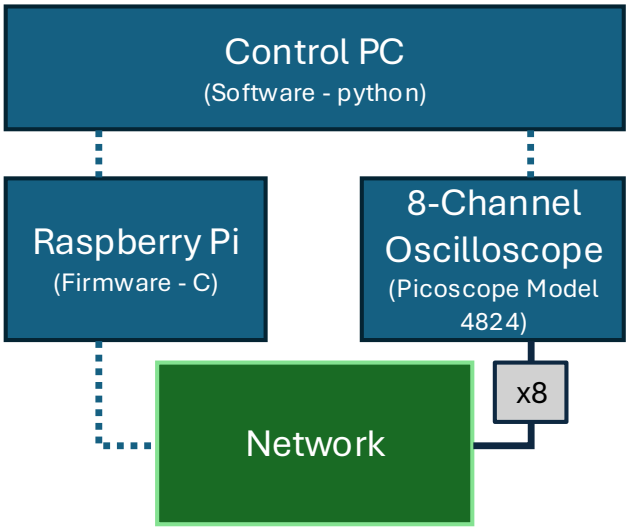

Figure E.3: Control and Readout. Diagram of the control and data acquisition system. The network is digitally controlled using SPI signals from a Raspberry Pi single board computer running the"firmware", which is written in C. Each node in the network is read out using an 8-channel, USB-connected oscilloscope. These subsystems are controlled from a main computer running custom control and acquisition software, written in Python. Dashed lines represent digital signals, solid lines represent analog signals.

E.2 Network Control

A schematic of the entire experimental system is shown in Fig. E.3. A PC orchestrates experiments through a custom Python codebase which communicates with each component of the system. First, the PC communicates with a Raspberry Pi single board computer which is used to generate the digital control signals which reprogram the network. The Raspberry Pi runs custom firmware which is written in C and built upon the bcm2835 library. This approach enables fast and precise programming of the network without latency from the PC. On the other side, the full analog output of each node in the network is captured by an 8-channel oscilloscope where data can either be stored or streamed to the PC. The signals can be captured in full fidelity with sampling rates as high as 20 MS/s or downsampled at rates as low as 100 kS/s. Even with downsampling, the amplitude and phase dynamics can be readily extracted.

All ICs in the network can be quickly reprogrammed in a time window much faster than the nodal amplitude relaxation time of 400-500 µs. In these experiments, we primarily focus on topology changes which are implemented solely by reprogramming the switch matrix. This is done by sending 240 bits to the four switch matrix

ICs via SPI. As the hardware is instantiated here, the switch matrices are updated asynchronously, meaning that the coupling turns on in four stages over the course of <80 µs. Notably, this is significantly faster than the timescale of the nodal relaxation dynamics. However, the next hardware generation implements a daisy chain configuration with a common latch which enables the switch matrices to update synchronously with a 50 ns nominal switch time. With this change, the topology can be fully reprogrammed in <50 µs and synchronously updated across the entire network. Faster programming speed can be achieved by optimizing the digital signal routing to enable faster clock rates.

E.3 Experimental Network Parameters

For the data presented in this thesis, the specific device parameters are presented in Table E.1. f_0 is the natural frequency of the device with no DC bias. The quality factor of the device is measured at a DC bias value which is close to the bias value where the experiments are conducted. The DC tunability is measured by applying voltage within ±4 V to the device and making a linear fit to the frequencies.

Table E.1: Properties of the NEMS Resonators

#	f_0 (MHz)	Q	DC Tunability (Hz/V)
1	1.799	2554	5204
2	1.792	2470	5344
3	1.805	2455	5336
4	1.810	2391	5296
5	1.822	2325	5232
6	1.807	2266	4768
7	1.799	2344	5256
8	1.806	2329	5332

Appendix F

ADDITIONAL EXPERIMENTAL METHODS

F.1 The Hilbert Transform

The Hilbert transform (HT) is a signal processing operation which calculates the analytic signal from a real-valued signal. It is primarily used to extract the instantaneous amplitude and phase of a narrow-band, nearly sinusoidal signal. Theoretically, to compute the Hilbert transform of a signal, you take the Fourier transform (FT) of a real-valued signal, $x(t)$; then, on negative frequency components, you shift the phase by $+\pi/2$ and on positive frequency components by $-\pi/2$. Taking the inverse Fourier transform of that spectrum produces the complex-valued analytic signal, $z(t)$. An approximation for the instantaneous amplitude of the signal is given by $a(t) = \text{abs}(z)$ and the instantaneous phase by $arg(z)$.

To extract amplitude and phase from the time-domain oscillation data, we use the Python package scipy [121] which has a built-in Hilbert transform function. Its algorithm works by taking a DFT of the signal, then it zeros the negative frequency components and doubles the positive frequency components. Then, the inverse DFT gives the analytic signal, from which the amplitude and phase can be extracted.

As a toy example, consider the signal $x(t) = A\sin(\omega_1 t + \phi)$. We want to extract A and ϕ. Rewriting the signal as

$$A\sin(\omega t + \phi) = -i\frac{1}{2}A\left[e^{i\phi}e^{i\omega t} - e^{-i\phi}e^{-i\omega t}\right]. \tag{F.1}$$

Recall that the Fourier transform of $e^{i\omega_1 t}$ is $\mathcal{F}(\omega) = 2\pi\delta(\omega - \omega_1)$. Therefore the FT of the function is

$$\mathcal{F}(\omega) = -i\frac{2\pi A}{2}\left[e^{i\phi}\delta(\omega - \omega_1) - e^{-i\phi}\delta(\omega + \omega_1)\right] \tag{F.2}$$

setting the negative frequency component to zero and doubling the positive frequency component gives

$$\mathcal{F}'(\omega) = -i2\pi A\left[e^{i\phi}\delta(\omega - \omega_1)\right]. \tag{F.3}$$

Taking the inverse FT, gives the analytic signal

$$z(t) = -iAe^{i\phi}e^{i\omega t} = -A\sin(\omega_1 t + \phi) + iA\cos(\omega_1 t + \phi). \tag{F.4}$$

Thus, we can recover the amplitude and phase as

$$|z(t)| = \sqrt{A^2\sin^2(\omega_1 t + \phi) + A^2\cos^2(\omega_1 t + \phi)} = A \tag{F.5}$$

and

$$\arg(z) = -\frac{\pi}{2} + \omega_1 t + \phi. \tag{F.6}$$

There are other methods that could be used to extract phase and amplitude. Using a PLL to track each oscillator in the network could reduce uncertainty in the phase extraction. Other methods come from state space estimation and include Kalman filter or maximum likelihood estimation. However, the Hilbert transform was found to be suitable for the work in this thesis.

F.2 Mutual Information

Mutual information (MI) is a pairwise measure of the amount of information shared between two random variables X and Y. In this work, we use mutual information to understand the structure of complex dynamical states of the oscillator network. We construct mutual information matrices by estimating the mutual information, I, between the phase trajectories of two oscillators as

$$I(X;Y) = H(X) + H(Y) - H(X,Y) \tag{F.7}$$

where $H(X)$ is the marginal Shannon entropy in bits, given by

$$H(X) = -\sum_{x=1}^{N_x} p_x \log_2(p_x) \tag{F.8}$$

and where $H(X,Y)$ is the joint Shannon entropy, given by

$$H(X,Y) = -\sum_{x=1}^{N_x}\sum_{y=1}^{N_y} p_{xy} \log_2(p_{xy}) \tag{F.9}$$

where N_x, N_y are the number of bins used for each random variable (we always take $N_x = N_y$), p_x is the probability of being in bin x, and p_{xy} is the probability of being in bin xy. In practice, we divide the phase space into $N_x \times N_y$ bins on the domain $X, Y \in (0, 2\pi]$ and compute these metrics using standard Python functions.

F.3 Calculating Allan Deviation

The Allan deviation is a measure of the frequency fluctuations of an oscillator over different averaging lengths. In comparison to phase noise measurements, it is useful for characterizing frequency fluctuations of oscillators on long timescales. We compute it by dividing a time series of frequency measurements, $f(t)$, into a set of N neighboring segments of length, τ. Letting $\bar{f}_i$ be the average frequency of segment i, the Allan variance as a function of τ is given by

$$\sigma_f^2(\tau) = \frac{1}{2N} \sum_{i=1}^{N-1} (\bar{f}_{i+1} - \bar{f}_i)^2 \tag{F.10}$$

where the Allan deviation is $\sigma_f(\tau) = \sqrt{\sigma_f^2(\tau)}$. The lowest possible τ is set by the time constant of the frequency measurement scheme, which is often a PLL. At low averaging times, we expect to see a rise in the Allan deviation by the power law $\sim \tau^{-1/2}$ resulting from white noise. The absolute minimum is set by the pink $(1/f)$ noise with power law $\sim \tau^0$, and at large averaging times we expect to see an increase due to drift which follows a $\sim \tau^1$ trend.

The Allan deviation is measured in Hz. The fractional Allan deviation is given by σ_f / f_0 where f_0 is the average frequency of the oscillator. The fractional Allan deviation is typically used in frequency shift mass spectroscopy applications because it can be used to directly approximate the mass resolution given the device's modal mass. In this work, we report the Allan deviation normalized by the resonator's linewidth so that they are direct measurements of parameter error in the theoretical model, where all parameters are also normalized to the linewidth.

BIBLIOGRAPHY

1. Strogatz, S. H. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering* (Westview Press, 1994).
2. Pikovsky, A., Rosenblum, M. & Kurths, J. *Synchronization* (Cambridge University Press, 2001).
3. Lu, C., Kim, M., Yang, Y., Gui, Y. & Hu, C.-M. Synchronization of dissipatively coupled oscillators. *Journal of Applied Physics* **134** (2023).
4. Salova, A. & D'Souza, R. M. Decoupled synchronized states in networks of linearly coupled limit cycle oscillators. *Physical Review Research* **2** (2020).
5. Bi, H. *et al.* Explosive oscillation death in coupled Stuart–Landau oscillators. *Europhysics Letters* **108,** 50003 (2014).
6. Perlikowski, P., Yanchuk, S., Popovych, O. & Tass, P. Periodic patterns in a ring of delay-coupled oscillators. *Physical Review E* **82,** 036208 (2010).
7. Nakao, H. Phase reduction approach to synchronisation of nonlinear oscillators. *Contemporary Physics* **57,** 188–214 (2016).
8. Kuramoto, Y. Self-entrainment of a population of coupled non-linear oscillators. *Lecture Notes in Physics* **39** (ed Araki, H.) 420–422 (1975).
9. Strogatz, S. H. Exploring complex networks. *Nature* **410,** 268–276 (2001).
10. Arenas, A., Díaz-Guilera, A., Kurths, J., Moreno, Y. & Zhou, C. Synchronization in complex networks. *Physics Reports* **469,** 93–153 (2008).
11. Van Hemmen, J. L. & Wreszinski, W. F. Lyapunov function for the Kuramoto model of nonlinearly coupled oscillators. *Journal of Statistical Physics* **72,** 145–166 (1993).
12. Wiley, D. A., Strogatz, S. H. & Girvan, M. The size of the sync basin. *Chaos* **16,** 015103 (2006).
13. Lu, J. & Steinerberger, S. Synchronization of Kuramoto oscillators in dense networks. *Nonlinearity* **33,** 5905 (2020).
14. Townsend, A., Stillman, M. & Strogatz, S. H. Dense networks that do not synchronize and sparse ones that do. *Chaos* **30,** 083142 (2020).
15. Kassabov, M., Strogatz, S. H. & Townsend, A. Sufficiently dense Kuramoto networks are globally synchronizing. *Chaos* **31** (2021).
16. Kassabov, M., Strogatz, S. H. & Townsend, A. A global synchronization theorem for oscillators on a random graph. *Chaos* **32** (2022).
17. Strogatz, S. H. From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena* **143,** 1–20 (2000).

18. Gómez-Gardeñes, J., Gómez, S., Arenas, A. & Moreno, Y. Explosive synchronization transitions in scale-free networks. *Physical Review Letters* **106,** 128701 (2011).

19. Leyva, I. *et al.* Explosive first-order transition to synchrony in networked chaotic oscillators. *Physical Review Letters* **108,** 168702 (2012).

20. Zhang, X., Hu, X., Kurths, J. & Liu, Z. Explosive synchronization in a general complex network. *Physical Review E* **88,** 010802 (2013).

21. D'Souza, R. M., Gómez-Gardenes, J., Nagler, J. & Arenas, A. Explosive phenomena in complex networks. *Advances in Physics* **68,** 123–223 (2019).

22. Kumar, M. & Gupta, S. Route to synchronization in coupled phase oscillators with frequency-dependent coupling: Explosive or continuous? *Physical Review E* **106,** 044310 (2022).

23. Iacopini, I., Petri, G., Barrat, A. & Latora, V. Simplicial models of social contagion. *Nature Communications* **10,** 2485 (2019).

24. Choi, J. & Kim, P. Critical neuromorphic computing based on explosive synchronization. *Chaos* **29,** 043110 (2019).

25. Kuehn, C. & Bick, C. A universal route to explosive phenomena. *Science Advances* **7,** eabe3824 (2021).

26. Sakaguchi, H. & Kuramoto, Y. A soluble active rotater model showing phase transitions via mutual entertainment. *Progress of Theoretical Physics* **76,** 576–581 (1986).

27. Fon, W. *et al.* Complex dynamical networks constructed with fully controllable nonlinear nanomechanical oscillators. *Nano Letters* **17,** 5977–5983 (2017).

28. Matheny, M. H. *et al.* Exotic states in a simple network of nanoelectromechanical oscillators. *Science* **363,** eaav7932 (2019).

29. Lifshitz, R. & Cross, M. C. Nonlinear dynamics of nanomechanical and micromechanical resonators. *Reviews of Nonlinear Dynamics and Complexity* **1,** 1–52 (2008).

30. Hoppensteadt, F. C. & Izhikevich, E. M. Synchronization of MEMS resonators and mechanical neurocomputing. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* **48,** 133–138 (2002).

31. Hoppensteadt, F. C. & Izhikevich, E. M. Synchronization of laser oscillators, associative memory, and optical neurocomputing. *Physical Review E* **62,** 4010 (2000).

32. Hoppensteadt, F. C. & Izhikevich, E. M. Oscillatory neurocomputers with dynamic connectivity. *Physical Review Letters* **82,** 2983 (1999).

33. Hoppensteadt, F. C. & Izhikevich, E. M. Pattern recognition via synchronization in phase-locked loop neural networks. *IEEE Transactions on Neural Networks* **11,** 734–738 (2000).

34. Shibata, T., Zhang, R., Levitan, S. P., Nikonov, D. E. & Bourianoff, G. I. *CMOS supporting circuitries for nano-oscillator-based associative memories* in *2012 13th International Workshop on Cellular Nanoscale Networks and their Applications* (2012), 1–5.

35. Levitan, S. P. *et al. Non-boolean associative architectures based on nano-oscillators* in *2012 13th International Workshop on Cellular Nanoscale Networks and their Applications* (2012), 1–6.

36. Nikonov, D. E. *et al.* Coupled-oscillator associative memory array operation for pattern recognition. *IEEE Journal on Exploratory Solid-State Computational Devices and Circuits* **1,** 85–93 (2015).

37. Frady, E. P. & Sommer, F. T. Robust computation with rhythmic spike patterns. *Proceedings of the National Academy of Sciences* **116,** 18050–18059 (2019).

38. Izhikevich, E. M. Weakly pulse-coupled oscillators, FM interactions, synchronization, and oscillatory associative memory. *IEEE Transactions on Neural Networks* **10,** 508–526 (1999).

39. Nagerl, J. & Berloff, N. G. Higher-order kuramoto oscillator network for dense associative memory. *arXiv preprint arXiv:2507.21984* (2025).

40. Hoppensteadt, F. C. & Izhikevich, E. M. *Associative memory of weakly connected oscillators* in *Proceedings of International Conference on Neural Networks (ICNN'97)* **2** (1997), 1135–1138.

41. Kumar, A. & Mohanty, P. Autoassociative memory and pattern recognition in micromechanical oscillator network. *Scientific Reports* **7,** 411 (2017).

42. Aoyagi, T. Network of neural oscillators for retrieving phase information. *Physical Review Letters* **74,** 4075 (1995).

43. Maffezzoni, P., Bahr, B., Zhang, Z. & Daniel, L. Oscillator array models for associative memory and pattern recognition. *IEEE Transactions on Circuits and Systems I: Regular Papers* **62,** 1591–1598 (2015).

44. Roychowdhury, J. Boolean computation using self-sustaining nonlinear oscillators. *Proceedings of the IEEE* **103,** 1958–1969 (2015).

45. Bösch, C., Roeder, G., Serra-Garcia, M. & Adams, R. P. Local learning rules for out-of-equilibrium physical generative models. *arXiv preprint arXiv:2506.19136* (2025).

46. Momeni, A. *et al.* Training of physical neural networks. *Nature* **645,** 53–61 (2025).

47. Csaba, G. & Porod, W. Coupled oscillators for computing: A review and perspective. *Applied Physics Reviews* **7,** 011302 (2020).

48. Theis, T. N. & Wong, H.-S. P. The end of moore's law: A new beginning for information technology. *Computing in Science & Engineering* **19,** 41–50 (2017).

49. Leiserson, C. E. *et al.* There's plenty of room at the Top: What will drive computer performance after Moore's law? *Science* **368,** eaam9744 (2020).

50. Leon, M. The escalating AI's energy demands and the imperative need for sustainable solutions. *WSEAS Transactions on Systems* **23,** 444–457 (2024).

51. Lin, L., Liao, X., Jin, H. & Li, P. Computation offloading toward edge computing. *Proceedings of the IEEE* **107,** 1584–1607 (2019).

52. McMahon, P. L. The physics of optical computing. *Nature Reviews Physics* **5,** 717–734 (2023).

53. Hari, R., Parkkonen, L. & Nangini, C. The brain in time: Insights from neuromagnetic recordings. *Annals of the New York Academy of Sciences* **1191,** 89–109 (2010).

54. Sporns, O., Tononi, G. & Kötter, R. The human connectome: A structural description of the human brain. *PLoS Computational Biology* **1,** e42 (2005).

55. Lucas, A. Ising formulations of many NP problems. *Frontiers in Physics* **2,** 5 (2014).

56. Honjo, T. *et al.* 100,000-spin coherent Ising machine. *Science Advances* **7,** eabh0952 (2021).

57. Inagaki, T. *et al.* A coherent Ising machine for 2000-node optimization problems. *Science* **354,** 603–606 (2016).

58. McMahon, P. L. *et al.* A fully programmable 100-spin coherent Ising machine with all-to-all connections. *Science* **354,** 614–617 (2016).

59. Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473,** 194–198 (2011).

60. Harrigan, M. P. *et al.* Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics* **17,** 332–336 (2021).

61. Yamaoka, M. *et al.* A 20k-spin Ising chip to solve combinatorial optimization problems with CMOS annealing. *IEEE Journal of Solid-State Circuits* **51,** 303–309 (2015).

62. Lo, H., Moy, W., Yu, H., Sapatnekar, S. & Kim, C. H. An Ising solver chip based on coupled ring oscillators with a 48-node all-to-all connected array architecture. *Nature Electronics* **6,** 771–778 (2023).

63. Chou, J., Bramhavar, S., Ghosh, S. & Herzog, W. Analog coupled oscillator based weighted ising machine. *Scientific Reports* **9,** 14786 (2019).

64. Wang, T., Wu, L., Nobel, P. & Roychowdhury, J. Solving combinatorial optimisation problems using oscillator based Ising machines. *Natural Computing* **20,** 287–306 (2021).

65. Mallick, A. *et al.* Using synchronized oscillators to compute the maximum independent set. *Nature Communications* **11,** 4689 (2020).

66. Berloff, N. G. *et al.* Realizing the classical XY Hamiltonian in polariton simulators. *Nature Materials* **16,** 1120–1126 (2017).

67. Kalinin, K. P. & Berloff, N. G. Simulating ising and n-state planar Potts models and external fields with nonequilibrium condensates. *Physical Review Letters* **121,** 235302 (2018).

68. Kalinin, K. P. & Berloff, N. G. Global optimization of spin Hamiltonians with gain-dissipative systems. *Scientific Reports* **8,** 17791 (2018).

69. Mohseni, N., McMahon, P. L. & Byrnes, T. Ising machines as hardware solvers of combinatorial optimization problems. *Nature Reviews Physics* **4,** 363–379 (2022).

70. Cucchi, M., Abreu, S., Ciccone, G., Brunner, D. & Kleemann, H. Hands-on reservoir computing: a tutorial for practical implementation. *Neuromorphic Computing and Engineering* **2,** 032002 (2022).

71. Teuscher, C. Revisiting the edge of chaos: Again? *Biosystems* **218,** 104693 (2022).

72. Carroll, T. L. Do reservoir computers work best at the edge of chaos? *Chaos* **30** (2020).

73. Nakajima, K. & Fischer, I. *Reservoir computing* (Springer, 2021).

74. Tanaka, G. *et al.* Recent advances in physical reservoir computing: A review. *Neural Networks* **115,** 100–123 (2019).

75. Bollt, E. On explaining the surprising success of reservoir computing forecaster of chaos? The universal machine learning dynamical system with contrast to VAR and DMD. *Chaos* **31,** 013108 (2021).

76. Pathak, J., Hunt, B., Girvan, M., Lu, Z. & Ott, E. Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical Review Letters* **120,** 024102 (2018).

77. Antonik, P., Gulina, M., Pauwels, J. & Massar, S. Using a reservoir computer to learn chaotic attractors, with applications to chaos synchronization and cryptography. *Physical Review E* **98,** 012215 (2018).

78. Wang, L., Fan, H., Xiao, J., Lan, Y. & Wang, X. Criticality in reservoir computer of coupled phase oscillators. *Physical Review E* **105,** L052201 (2022).

79. Velichko, A., Ryabokon, D., Khanin, S., Sidorenko, A. & Rikkiev, A. *Reservoir computing using high order synchronization of coupled oscillators* in *IOP Conference Series: Materials Science and Engineering* **862** (2020), 052062.

80. Yamane, T., Katayama, Y., Nakane, R., Tanaka, G. & Nakano, D. *Wave-based reservoir computing by synchronization of coupled oscillators* in *International Conference on Neural Information Processing* (2015), 198–205.

81. Kanao, T. *et al.* Reservoir computing on spin-torque oscillator array. *Physical Review Applied* **12,** 024052 (2019).

82. Dion, G., Mejaouri, S. & Sylvestre, J. Reservoir computing with a single delay-coupled non-linear mechanical oscillator. *Journal of Applied Physics* **124** (2018).

83. Kartal, E. *et al.* Nanomechanical Systems for reservoir computing applications. *Advanced Intelligent Systems,* 2400971 (2025).

84. Sun, J. *et al.* Novel nondelay-based reservoir computing with a single micromechanical nonlinear resonator for high-efficiency information processing. *Microsystems & Nanoengineering* **7,** 83 (2021).

85. Van der Sande, G., Brunner, D. & Soriano, M. C. Advances in photonic reservoir computing. *Nanophotonics* **6,** 561–576 (2017).

86. Vettelschoss, B., Röhm, A. & Soriano, M. C. Information processing capacity of a single-node reservoir computer: an experimental evaluation. *IEEE Transactions on Neural Networks and Learning Systems* **33,** 2714–2725 (2021).

87. Mandal, S., Sinha, S. & Shrimali, M. D. Machine-learning potential of a single pendulum. *Physical Review E* **105,** 054203 (2022).

88. Alomar, M. L. *et al.* Digital implementation of a single dynamical node reservoir computer. *IEEE Transactions on Circuits and Systems II: Express Briefs* **62,** 977–981 (2015).

89. Appeltant, L. *et al.* Information processing using a single dynamical node as complex system. *Nature Communications* **2,** 468 (2011).

90. Asano, M., Okamoto, H. & Yamaguchi, H. Synthesized Kuramoto potential via optomechanical Floquet engineering. *Science Advances* **11,** eady4167 (2025).

91. Cleland, A. N. *Foundations of Nanomechanics* ISBN: 978-3-642-07821-7 (Springer Berlin Heidelberg, 2003).

92. *Piezoelectric MEMS Resonators* (eds Bhugra, H. & Piazza, G.) ISBN: 978-3-319-28686-0 (Springer International Publishing, 2017).

93. Schmid, S., Villanueva, L. G. & Roukes, M. L. *Fundamentals of Nanomechanical Resonators* (Springer, 2016).

94. Postma, H., Kozinsky, I., Husain, A. & Roukes, M. Dynamic range of nanotube-and nanowire-based electromechanical systems. *Applied Physics Letters* **86** (2005).

95. Nayfeh, A. H. & Mook, D. T. *Nonlinear Oscillations* (John Wiley & Sons, 1979).

96. Lu, Y., Shao, Q., Amabili, M., Yue, H. & Guo, H. Nonlinear vibration control effects of membrane structures with in-plane PVDF actuators: A parametric study. *International Journal of Non-Linear Mechanics* **122,** 103466 (2020).

97. Zaghloul, U. & Piazza, G. Highly scalable NEMS relays with stress-tuned switching voltage using piezoelectric buckling actuators. *IEEE Transactions on Electron Devices* **61,** 3520–3528 (2014).

98. Erbil, S. O. *et al.* Full electrostatic control of nanomechanical buckling. *Physical Review Letters* **124,** 046101 (2020).

99. Piazza, G., Felmetsger, V., Muralt, P., Olsson III, R. H. & Ruby, R. Piezoelectric aluminum nitride thin films for microelectromechanical systems. en. *MRS Bulletin* **37,** 1051–1061 (2012).

100. Karabalin, R. B. *et al.* Piezoelectric nanoelectromechanical resonators based on aluminum nitride thin films. *Applied Physics Letters* **95,** 103111 (2009).

101. Karabalin, R. B., Villanueva, L. G., Matheny, M. H., Sader, J. E. & Roukes, M. L. Stress-induced variations in the stiffness of micro- and nanocantilever beams. *Physical Review Letters* **108,** 236101 (2012).

102. Matheny, M. *Nonisochronous oscillations in piezoelectric nanomechanical resonators* PhD thesis (California Institute of Technology, 2012).

103. Matheny, M. H., Villanueva, L. G., Karabalin, R. B., Sader, J. E. & Roukes, M. L. Nonlinear mode-coupling in nanomechanical systems. *Nano Letters* **13,** 1622–1626 (2013).

104. Pobedinskas, P. *et al.* Thickness dependent residual stress in sputtered AlN thin films. *Thin Solid Films* **522,** 180–185 (2012).

105. Aubin, K. *et al.* Limit cycle oscillations in CW laser-driven NEMS. *Journal of Microelectromechanical Systems* **13,** 1018–1026 (2004).

106. Houri, S., Asano, M., Okamoto, H. & Yamaguchi, H. Self-sustained libration regime in nonlinear microelectromechanical devices. *Physical Review Applied* **16,** 064015 (2021).

107. Karabalin, R. B., Cross, M. C. & Roukes, M. L. Nonlinear dynamics and chaos in two coupled nanomechanical resonators. *Physical Review B* **79** (2009).

108. Colombano, M. *et al.* Synchronization of optomechanical nanobeams by mechanical interaction. *Physical Review Letters* **123,** 017402 (2019).

109. Bhaskar, A., Walth, M., Rand, R. H. & Zehnder, A. T. Synchronization in pairs of opto-thermally driven mechanically coupled micro-oscillators. *Journal of Microelectromechanical Systems* **32,** 136–142 (2023).

110. Agrawal, D. K., Woodhouse, J. & Seshia, A. A. Observation of locked phase dynamics and enhanced frequency stability in synchronized micromechanical oscillators. *Physical Review Letters* **111,** 084101 (2013).

111. Agrawal, D. K., Woodhouse, J. & Seshia, A. A. Synchronization in a coupled architecture of microelectromechanical oscillators. *Journal of Applied Physics* **115,** 164904 (2014).

112. Zehnder, A. T., Rand, R. H. & Krylov, S. Locking of electrostatically coupled thermo-optically driven MEMS limit cycle oscillators. *International Journal of Non-Linear Mechanics* **102,** 92–100 (2018).

113. Pu, D., Wei, X., Xu, L., Jiang, Z. & Huan, R. Synchronization of electrically coupled micromechanical oscillators with a frequency ratio of 3:1. *Applied Physics Letters* **112,** 013503 (2018).

114. Zhang, M. *et al.* Synchronization of micromechanical oscillators using light. *Physical Review Letters* **109,** 233906 (2012).

115. Li, J. *et al.* All-optical synchronization of remote optomechanical systems. *Physical Review Letters* **129,** 063605 (2022).

116. Erokhin, S. & Berkov, D. Robust synchronization of an arbitrary number of spin-torque-driven vortex nano-oscillators. *Physical Review B* **89,** 144421 (2014).

117. Kenig, E. *et al.* Optimal operating points of oscillators using nonlinear resonators. *Physical Review E* **86,** 056207 (2012).

118. Villanueva, L. G. *et al.* Surpassing fundamental limits of oscillators using nonlinear resonators. *Physical Review Letters* **110,** 177208 (2013).

119. Matheny, M. H. *et al.* Phase synchronization of two anharmonic nanomechanical oscillators. *Physical Review Letters* **112,** 014101 (2014).

120. Bahr, B. W., Popa, L. C. & Weinstein, D. *16.8 1GHz GaN-MMIC monolithically integrated MEMS-based oscillators* in *2015 IEEE International Solid-State Circuits Conference-(ISSCC) Digest of Technical Papers* (2015), 1–3.

121. Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17,** 261–272 (2020).