

**Enhancing the scan rate for axion dark matter:
Quantum noise evasion and maximally informative analysis**

by

Daniel A. Palken

B.A., Bowdoin College, 2014

M.S., University of Colorado, 2017

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2020

This thesis entitled:
Enhancing the scan rate for axion dark matter:
Quantum noise evasion and maximally informative analysis
written by Daniel A. Palken
has been approved for the Department of Physics

Konrad Lehnert

Eric Cornell

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Palken, Daniel A. (Ph.D., Physics)

Enhancing the scan rate for axion dark matter:

Quantum noise evasion and maximally informative analysis

Thesis directed by Prof. Konrad Lehnert

Dark matter axions, particles whose hypothesized existence could resolve two of the largest outstanding mysteries in physics, are made difficult to detect by their extremely feeble coupling to ordinary matter. The axion haloscope, first realized experimentally three decades ago, remains among the most viable detection platforms, but even today's leading technology would optimistically require many millennia to scan commonly targeted portions of axion parameter space. Today's axion direct detection community is searching therefore not just for the elusive particles, but for technologies and innovations that will permit more efficient searches. In this thesis, I present two such innovations. First, the resource of quantum squeezing sequesters the noise of a haloscope measurement into unmeasured observables, improving the sensitivity bandwidth and hence the scan rate of the detector, its primary figure of merit. Second, valuable information pertinent to the existence of the axion is discarded by the standard hypothesis testing framework used to look for axions. An alternative, Bayesian analysis utilizes the full information content of the haloscope measurement, yielding a tangible speedup at zero operational or hardware cost. Together, these innovations are used improve threefold the scan rate of a dark matter search performed by the Haloscope At Yale Sensitive To Axion Cold dark matter (HAYSTAC) experiment.

Dedication

In loving memory of Marcia Linsey

Acknowledgements

It is no overstatement to say that none of the work in this thesis would have been possible without the support, mentorship, collaboration, and friendship of the many wonderful people who have come into my life since moving to Colorado, or were already a fixture in it beforehand.

The foundation upon which my ability to get up each morning and work towards something larger than myself is built are my family and friends. To my mother and father, Judy and Rich Palken, an infinity of thank yous for all you have done and continue to do to support my education (and to dad in particular: I am sorry I never built the warp-drive you wanted; we will have to find another way to teach the aliens about Moe, Larry, and Curly before they get to meet you). To my brother Greg: thank you for always reminding me of the joy of stepping outside for a breath of fresh air, and a quick spin up a series of mountain passes. To Peyton, Sam, James, and my other non-Colorado friends who have been a part of my life long before graduate school and will continue to be so long after it: I appreciate the asymmetry in who has flown out to visit whom these past several years, and the subdued chuckle (or maybe two) we've shared on such occasions. To Vivian, Ursula, and Melinda, I am delighted our paths intersected for a time in a second location, and to John, I look forward to a third. To the many close friends I have made during my time in Colorado — Adam, Bryce, Karl, Kelsey, Kristina, Lindsay, Marissa, and Steven, to name a subset who are not singled out below — thank you for all the adventures, board games, meals, and conversations. To the many volunteers I have worked with and mentors I have had at the Citizens' Climate Lobby, thank you for all you have done for me, and continue to do every day for others.

It has now been nearly a decade since I first took the introductory astronomy course with

Professor Thomas Baumgarte at Bowdoin College with which set me on the path of pursuing a physics degree. That ambition could well have been derailed by my engrossment, that very same semester, in an English literature course modernity taught by Professor Ann Kibbie. Instead, my professors at Bowdoin encouraged me to pursue both, and the humanities taught me a deep appreciation for the sciences which is among the greatest gifts I've ever received. In addition to that, my other physics professors at Bowdoin, Dale Syphers, Mark Battle, Stephen Naculich, and my undergraduate advisor Madeleine Msall, taught me more than a bit of physics as well. More importantly, they taught me how to learn what I did not already know.

The physics I have learned since, and the experiments I have lent a hand to in my time at CU Boulder and JILA, featured a supporting cast layers deep which constantly lifted me up and put me squarely in a position where I could not but succeed. It has been a joy to work with the talented and scrupulous members of the HAYSTAC collaboration towards the goal of detecting dark matter. Within JILA, The Lehnert Lab made for an environment where quantum circuits were fun, and so was eating lunch together most days at the C4C. Will Kindel, Adam Reed, Jeremie Viennot, Ben Chapman, Pete Burns, Xizheng Ma, Lucas Sletten, Eric Rosenthal, Max Urmey, and Rob Delaney featured especially prominently throughout my time in there, and you all have my sincere thanks.

In addition to colleagues from my own generation, I have had the opportunity to work with and be informally advised by one-of-a-kind physicists who have been at it for considerably longer (though you would know as much only from speaking to, and certainly not looking at them). Reina Maruyama brings a broad knowledge of dark matter detection to our team. Karl van Bibber knows much about dark matter axions, which we both very much hope exist, but he is all the while an existence-proof that optimism and good spirit should be thought of as properties of the journey, destination come when it will. Eric Cornell absorbs new information with a coupling rate suggestive of a 2π error, and it was one of the great delights of my time in JILA to get to work with him on an unforeseen project of unforeseen scope, all for unseen dark matter. In addition, I extend my further thanks to Eric, as well as to Manny Knill, Kevin Stenson, and Vanja Dukic, for taking the time to read and engage with my work as members of my thesis committee.

The core group of colleagues I collaborated with include Tanya Roussy of the Cornell lab, who proved a kindred spirit in ruthlessly applying Bayes' theorem to look for the missing parts of our universe, and Kelly Wurtz, who arrived in the Lehnert Lab knowing nothing about quantum circuits, and left less than two years later having taught me how they could be used to improve tenfold upon one of the two central ideas of this thesis.

And then there are Kelly Backes, Max Malnou, and Ben Brubaker. Winning the lottery turns out to be underrated; I can attest to this, having done so now on three occasions.

Max Malnou joined the Lehnert Lab at a time when I was still finding my place in both group and collaboration. I had a passion for learning from books, but struggled to make the connection to practical application. I credit Max with changing that. He proved an adept and patient mentor and colleague, and modeled what I strove to become as a physicist. By the end of his time at JILA, the two of us shared a brain on many matters, and knew how to complement the work style and strengths of each other on what remained. It was an absolute joy and privilege having a partner to collaborate so closely with on such fascinating projects.

The small group at Yale which runs the axion haloscope around which all the work in this thesis revolves could well have chosen any of a number of people to step into the very large shoes of their departing force of nature of a graduate student. I very much doubt any of those choices would have been better suited to the task at hand than Kelly Backes. Kelly involves herself in every detail, grand or minutial, of her experiment. For a project like ours with so many moving parts, this is as difficult as it is essential. Her diligence and eagerness are in equal measure contagious, and she is a joy to sit in an office with and bang one's head against a wall for hours on end — an activity known to cause headaches in most other contexts.

The shoes that Kelly now comfortably wears were previously filled by Ben Brubaker. Ben, where do I begin? We knew each other, though not terribly well, my first few years of graduate school: you were busy building from the ground up the experiment I would inherit a role on for my own research. It was when you came to JILA to work in the Lehnert Lab for your post-doctoral research that our more direct partnership began to flourish. Speaking very conservatively,

you've taught me the vast majority of what I know (and also some that I don't know) about dark matter, and you encouraged and facilitated a dive into statistics which has permanently changed my thinking in ways that extend well beyond the applications at hand. In a few short years, you've also read and reread more dubious grammatical constructions (including this one) that I've written into more scientific documents (including this one) than I can ever rightly thank you for. In doing so, you've helped me clarify the underlying thoughts for myself and others, and made me a better scientist as part of the bargain.

Lastly, to my advisor, Konrad Lehnert: your dedication to and stamina for scientific inquiry set a standard that everyone who works with you admires and attempts to emulate. I've seen you exhibit integrity when there was risk in doing so, and put the ambitions and goals of your students and post-docs ahead of your own, somehow without ever sacrificing on the quality of work that comes out of the lab. You have a keen sense of what matters and what does not, in both the physics world and the human one, and it all shows in how you choose to allocate your time. Based on the resources that have flowed to myself and others in the lab these past years, I am pretty sure you must be writing a lot of grants, but I would be hard pressed to say when (Wednesday mornings, perhaps?), given the seemingly uninterrupted meeting schedule you've constructed around all of the different physics and physicists you oversee. That simultaneous attention to detail and to people has made the Lehnert Lab a place where I have not only gotten to contribute to science, but have enjoyed doing so.

Contents

Chapter

1	Thesis structure	1
2	Dark matter, axions, and haloscopes	5
2.1	Evidence of dark matter	6
2.1.1	Galactic velocity dispersion in the Coma Cluster	6
2.1.2	Galactic rotation curves	7
2.1.3	Further arguments for and against dark matter	8
2.2	The not-so-miraculous WIMP	11
2.3	The QCD axion solution to dark matter	12
2.3.1	The strong CP problem	13
2.3.2	The Peccei-Quinn mechanism	14
2.3.3	Axions as dark matter candidates	14
2.4	Detectable features of axionic dark matter	17
2.4.1	Axion parameter space	17
2.4.2	Local axion dark matter density	19
2.4.3	Axion lineshape	19
2.5	The axion haloscope	21
2.5.1	Previous haloscope results	23
2.5.2	Haloscope scan rate	25
2.5.3	Improving the scan rate	26

3	Josephson parametric amplifiers	29
3.1	The Josephson junction	29
3.1.1	Josephson inductance	31
3.1.2	Josephson energy	32
3.2	The DC SQUID	32
3.2.1	Critical current and inductance	33
3.2.2	Geometrical inductance and flux-trapping	37
3.3	Circuit quantization	39
3.3.1	LC circuit	39
3.3.2	Lagrangian and Hamiltonian	40
3.3.3	Classical equations of motion	40
3.3.4	Quantization and ladder operators	41
3.3.5	Heisenberg equation of motion	42
3.4	Lumped-element Josephson parametric amplifier	42
3.4.1	Basic TKC circuit analysis	45
3.4.2	TKC Lagrangian and Hamiltonian	47
3.4.3	Kerr Hamiltonian	48
3.4.4	Input-output theory derivation of gain	50
3.4.5	Gain-bandwidth product	55
3.4.6	Single-quadrature amplification, deamplification, and squeezing	56
3.4.7	Derivation of flux-pumped JPA gain	58
3.5	Validity of the Duffing approximation	62
3.5.1	The Qp condition for a single-junction JPA	63
3.5.2	Significance of the participation ratio	65
3.5.3	From junction to array	65
3.5.4	Series RLC and general Thévenin case	66
3.5.5	Parallel RLC and general Norton case	68

3.5.6	The Qp condition for an array-based JPA	69
3.5.7	Independence of the validity of the Duffing approximation from the nonlinearity	70
3.5.8	Modified condition for the Duffing approximation	71
3.6	Current-pumped JPA implementation	73
3.6.1	SQUID array specifications	73
3.6.2	TKC specifications and simulations	74
3.6.3	Characterization measurements	77
3.7	Flux-pumped JPA implementation	81
3.7.1	Flux-pumped JPA design	83
3.7.2	Characterization measurements	90
4	Optimal current-pumped Josephson parametric amplifier operation for vacuum squeezing	95
4.1	Signal transformation in phase space	97
4.2	Coherent state deamplification	100
4.3	Optimizing vacuum squeezing with a current-pumped JPA	103
4.4	Experimental details and calibration	105
4.4.1	Experimental setup	106
4.4.2	Inter-JPA loss calibration	108
4.4.3	Added noise calibration	109
4.5	Concluding remarks on current-pumped JPAs optimized for squeezing	110
5	Squeezed vacuum used to accelerate the search for a weak classical tone	112
5.1	Theory of the squeezed state receiver	113
5.1.1	Lossless calculation	118
5.1.2	Calculation with transmission losses	123
5.1.3	Equivalence between single- and double-quadrature amplification	126
5.2	Model for the axion field	128
5.3	Enhanced signal visibility with squeezing	132

5.4	Mock-haloscope experiment: accelerated search for a weak classical tone	135
5.5	Experimental details and calibration	138
5.5.1	Experimental setup	138
5.5.2	Phase-locking of SQ and AMP pumps	140
5.5.3	Intra-SSR loss measurement	142
5.6	Complementary measurements of the microwave tone's improvement in resolution . .	143
5.7	Processing of the SSR spectra	145
5.8	Concluding remarks on the SSR for enhancing the scan rate of real haloscope searches	148
6	Axion haloscope processing and analysis	149
6.1	Haloscope processing overview	149
6.2	Axion search analysis considerations	154
6.2.1	Distinguishing features of axion and haloscope searches	156
6.3	Frequentist threshold framework	159
6.3.1	Framework details	159
6.3.2	Aggregate exclusion	162
6.4	p -value framework	164
6.5	Look-elsewhere effect: implications for detection and exclusion	167
6.5.1	Frequentist frameworks and the look-elsewhere effect	172
7	A maximally informative axion dark matter analysis	174
7.1	Bayesian power-measured framework	175
7.1.1	Updated belief	176
7.1.2	Aggregation and the look-elsewhere effect	178
7.1.3	Choice of priors	180
7.1.4	Correlations between grand spectrum bins	183
7.1.5	Practical departures from Gaussianity	184
7.2	Effects of data processing on information content	186

7.3	Comparison of analysis frameworks	190
7.3.1	Bayesian threshold 1 framework	190
7.3.2	Bayesian threshold 2 framework	192
7.3.3	Outcome distributions for Bayesian analyses	195
7.4	Test using HAYSTAC Phase I dataset	200
7.4.1	Main results	202
7.4.2	Least-unlikely candidates	203
7.4.3	Additional features of the Phase I dataset	204
7.5	Concluding remarks on statistical analysis frameworks for axion searches	205
8	A quantum-enhanced search for dark matter axions	207
8.1	Setup and operation	208
8.2	Data processing and analysis	216
8.2.1	Subspectra and raw spectra	218
8.2.2	Processed spectra	219
8.2.3	Rescaled spectra	222
8.2.4	Combined spectrum	223
8.2.5	Grand spectrum	225
8.3	Results	226
8.3.1	JILA analysis results	226
8.3.2	Official HAYSTAC Phase II results and conclusions	227
8.4	Calibration measurements	229
8.4.1	Squeezing and cavity spectral density measurements	231
8.4.2	Added noise measurement	232
9	Search for axion-like dark matter with the JILA eEDM experiment	234
9.1	ALP sensitivity of the JILA eEDM dataset	236
9.2	Spectral content of the JILA eEDM dataset	238

9.3	Quantifying the hypotheses	241
9.3.1	No-ALP hypothesis	241
9.3.2	ALP hypothesis	243
9.3.3	Coherent ALP	245
9.3.4	Partially coherent ALP	248
9.3.5	Finite sample time correction	251
9.4	ALP exclusion results	252
9.5	Prospects for future JILA eEDM ALP searches	254
10	Conclusion and outlook	256
10.1	Conclusion	257
10.2	Outlook	259
	Bibliography	264
	Appendix	
A	NIST trilayer fabrication process	290
B	Resolution to the flux-tuning problem for flux-pumped JPAs	293
C	Effect of improper tuning rod thermalization on scan rate	297
D	List of Abbreviations	301

Tables

Table

5.1	Representative physical values used to estimate axion port parameters	130
5.2	Model parameter values calculated using the physical values from Table 5.1	132
7.1	Sample set of probabilities used to estimate prior axion hypothesis probability	182
7.2	Comparison of statistical analysis frameworks	194
9.1	ALP coherence for high and low frequency ALP analyses and HAYSTAC Phase II	250
C.1	Theoretical parameters for the HAYSTAC Phase II haloscope with and without a hot rod and/or squeezing	299

Figures

Figure

2.1	Galactic rotation curves	8
2.2	Bullet Cluster collision	10
2.3	WIMP parameter space	12
2.4	Axion-photon parameter space prior to this thesis	18
2.5	Virialized axion lineshapes	22
2.6	Electric and magnetic fields in simple and tunable cylindrical cavities	24
2.7	Axion-photon exclusion obtained using haloscopes	25
3.1	Representations of a Josephson junction	31
3.2	DC SQUID circuit representation	34
3.3	Vector potential integration path for a SQUID	34
3.4	Total versus externally applied flux for a SQUID with different geometrical inductances	38
3.5	Simple LC circuit	39
3.6	Circuit schematic and IOT model for a TKC	44
3.7	Relationship between TKC and JPA	45
3.8	Duffing oscillator response	54
3.9	Relation between one- and two-quadrature gain	57
3.10	PNL and SNL	63
3.11	Simple embedding circuits	64

3.12 Series inductor algebra showing the equivalence between stray and array linear inductances	67
3.13 Parallel inductor algebra showing the equivalence between stray and array linear inductances	68
3.14 Duffing and higher-order oscillator responses at different quality factors	72
3.15 Josephson junction critical current measurements	75
3.16 Superconducting niobium wiring layout for a pair of SQUIDs	76
3.17 SEM images of a TKC for a current-pumped JPA	76
3.18 Layouts and design parameters for current-pumped TKCs	77
3.19 Simulated responses and parameter values for current-pumped JPAs	78
3.20 JPA configuration for current pumping	79
3.21 Resonance frequency measurements for current-pumped JPAs	80
3.22 Flux tuning of current-pumped JPA resonance with a DC current	81
3.23 Gain and bandwidth measurements for a current-pumped JPA	82
3.24 Discarded flux-pumped JPA designs	84
3.25 Final flux-pumped JPA design	85
3.26 EM simulations of a flux-pumped JPA	87
3.27 Flux-pumped JPA scattering parameter simulations	88
3.28 Flux-pumped JPA gain simulations	89
3.29 JPA configuration for flux pumping	89
3.30 SEM images of a TKC for a flux-pumped JPA	90
3.31 Photo of a mounted flux-pumped JPA	91
3.32 Flux tuning of flux-pumped JPA resonance with a DC current	92
3.33 Flux-pumped JPA gain curves showing tunability between and within different coil tunings	93
3.34 Flux-pumped JPA gains at multiple operating points	94
3.35 Power-handling measurements for flux- and current-pumped JPA operation	94

4.1	Theory of distortion in a JPA	98
4.2	Experimental setup for coherent tone distortion and vacuum squeezing	100
4.3	Dependence of coherent tone distortion upon amplifier operating point	101
4.4	Dependence of deamplification upon amplifier operating point	102
4.5	Vacuum squeezing with current-pumped JPAs	104
4.6	Dependence of vacuum squeezing on operating point	105
4.7	Full experimental setup used for both coherent tone deamplification and squeezing	107
4.8	VTS noise calibration	110
5.1	Schematic of the SSR and cavity and illustration of the benefit of squeezing	115
5.2	Theoretical scan rate enhancement for idealized and realistic haloscope models	118
5.3	Equivalent modelling of the fictitious axion port	130
5.4	Photos of the AI cavity used to test the SSR	133
5.5	Squeezing achieved in the SSR	134
5.6	Visibility improvement from squeezing	135
5.7	Data processing and results from many trials of the mock-haloscope experiment	137
5.8	Full experimental setup used for the mock-haloscope experiment	141
5.9	Intra-SSR transmission loss calibration measurement	143
5.10	Complementary measurements of signal visibility in the SSR	144
6.1	Gaussian probability distributions for no-axion and axion distributions	155
6.2	Illustrative grand spectrum data representation for standard haloscope subject to an FT analysis	160
6.3	Effective correction to the required SNR from the look-elsewhere effect	171
7.1	Four possible outcomes for a two-scan FT protocol	192
7.2	Outcome probability distributions for total and aggregate prior updates	197
7.3	HAYSTAC Phase I BPM reanalysis	201

8.1	Simplified HAYSTAC Phase II experimental schematic	209
8.2	HAYSTAC cavity photo	210
8.3	Photos of the LD250 BlueFors dilution refrigerator used for HAYSTAC Phase II	210
8.4	Field specifications for the solenoid magnet used in HAYSTAC	211
8.5	Mixing chamber stage layout and photo of magnetic shielding	212
8.6	Measurement of JPA tunability with magnetic field on and off	213
8.7	Squeezing measured in HAYSTAC Phase II	213
8.8	Illustration of the SSR-equipped haloscope showing the transformation of the vacuum state in quadrature space	214
8.9	Advantage conferred by squeezing calculated with parameters used in HAYSTAC	215
8.10	Raw spectrum cuts for HAYSTAC Phase II initial scan	217
8.11	HAYSTAC Phase II raw spectra	219
8.12	Individualized SG filtering	220
8.13	HAYSTAC Phase II processed spectra	221
8.14	Rescaling visibility and individual rescaled spectrum	222
8.15	HAYSTAC Phase II initial scan combined spectrum	224
8.16	HAYSTAC Phase II initial scan grand spectrum	225
8.17	Results of the JILA analysis of the HAYSTAC Phase II dataset	226
8.18	Final HAYSTAC Phase II axion exclusion	228
8.19	Model used for HAYSTAC Phase II noise calibration measurements	229
8.20	Sample calibration measurements to determine cavity spectral density and squeezer gain	232
9.1	JILA eEDM dataset	238
9.2	LSSA of the JILA eEDM dataset	239
9.3	Phase sensitivity of the JILA eEDM dataset	242
9.4	Construction of the no-ALP and ALP distributions in quadrature space	245

9.5	Effect of ALP field decoherence on the ALP-only distribution variance	250
9.6	Effect of finite sample time on the ALP-only distribution variance	251
9.7	ALP exclusion results	253
9.8	Astrophysical and laboratory limits on the ALP-gluon interaction	254
10.1	Axion namesake-inspired representation of the SSR	256
10.2	Impact on scan rate of decreasing haloscope noise over time	258
10.3	Concept for future quantum enhancement to haloscope scan rate	262
A.1	NIST trilayer fabrication process	292
B.1	Location of cuts made with a FIB on dummy SQUIDS	293
B.2	Bidirectional flux-tuning sweep of FIB-cut JPA showing fixed tuning problem	294
B.3	Gain map and traces for a FIB-cut JPA	295
C.1	Theoretical signal and noise spectral densities with and without hot rod and squeezing	298
C.2	Theoretical axion visibilities with and without hot rod and squeezing	299
C.3	Theoretical scan rates with and without hot rod and squeezing	300

Chapter 1

Thesis structure

We do not know what constitutes the cosmos, and it may take us a very long time to find out. The problem is that of dark matter, and the difficulty of directly detecting it. This thesis sets out to speed up the search for dark matter axions through the use of quantum measurement and statistical analysis.

I begin in Ch. 2 with a motivating discussion of the problem of dark matter, the axion as a potential solution, and the axion haloscope as the primary device used in efforts to detect it. In Sec. 2.1, I overview some of the most prominent evidence for dark matter. Section 2.2 discusses the weakly interacting massive particle (WIMP),¹ a once highly promising dark matter candidate now somewhat disfavored by experimental null results. Section 2.3 then turns to the candidate most central to this thesis, the axion of quantum chromodynamics. This section and the ensuing Sec. 2.4 discuss those properties that qualify axions as dark matter candidates, and those pertinent to detecting their signature experimentally, respectively. In Sec. 2.5, I introduce the axion haloscope. The chapter concludes with a quantitative discussion of what limits the performance of axion haloscopes, as parametrized by their scan rates — the speeds at which they can move through parameter space in search of dark matter axions.

Chapter 3 introduces Josephson parametric amplifiers (JPAs), devices that can create and process squeezed states of microwave light, which will ultimately enable me to enhance the scan rate of an axion haloscope. This chapter begins in Sec. 3.1 with a discussion of the Josephson

¹ A full list of acronyms and other abbreviations appearing in this thesis can be found in Appx. D.

junction, whose special properties will be used to enable the noiseless amplification and squeezing later in this thesis. I derive in Sec. 3.2 key properties of pairs of Josephson junctions known as superconducting quantum interference devices (SQUIDs). After Sec. 3.3 provides the language of circuit quantization, Sec. 3.4 combines the building blocks developed thus far into a JPA. I derive how JPAs achieve gain and squeezing, before turning in Sec. 3.5 to a discussion of the validity of the Duffing approximation, a key simplification made in the JPA gain derivation. The implications of this discussion are important for understanding the physics of array-based JPAs as a partial remedy to the problem of limited power handling. Finally, turning our attention to the devices themselves, Secs. 3.6 and 3.7 discuss the design and characterization of JPAs intended to be operated with different types of pump tones. The current-pumped JPAs of Sec. 3.6 will be central to Ch. 4, while the flux-pumped JPAs of Sec. 3.7 will feature in Chs. 5 and 8. The fabrication process, performed at NIST Boulder, is discussed in Appx. A, and the resolution to a flux-tuning problem that plagued our flux-pumped JPAs is discussed in Appx. B.

Chapter 4 presents a technique to improve the power handling of current-pumped JPAs used for squeezing. I begin in Sec. 4.1 with a theoretical treatment of how a JPA transforms a signal in phase space, drawing on results from Ch. 3. Section 4.2 details an experiment performed with coherent states of the microwave field whose amplitudes are comparable to vacuum fluctuations. The results of this experiment suggest that fragility to fluctuations in the power of a JPA's pump tone can be traded off for improved squeezing. In Sec. 4.3, I show in a second experiment performed on vacuum states that this result indeed improves the squeezing that can be delivered to a physical system such as an axion haloscope. Section 4.4 describes experimental details and calibration measurements for the transmission losses and added noise in the setup used throughout this section. The improvement achieved in this chapter, summarized in Sec. 4.5, is limited by presently achievable degrees of squeezing (in turn limited by transmission losses within the system), suggesting that the superior options for parametric amplification are the flux-pumped JPAs used in ensuing chapters.

Chapter 5 modifies the experimental architecture of Ch. 4 from one for testing squeezing, to one for using it: the integration of two JPAs into an axion haloscope-like setup forms a squeezed

state receiver (SSR) apparatus capable of exploiting the uncertainty principle of quantum mechanics to enhance the search rate for an axion-like signal. Section 5.1 uses input-output theory to show that the SSR works by eliciting a favorable tradeoff between peak haloscope sensitivity and bandwidth, and includes a detailed theoretical calculation with experimental parameters used throughout the chapter. Section 5.2 mathematically models how the axion field interacts with a haloscope. Sections 5.3 and 5.4 report on a pair of experiments to demonstrate the efficacy of the SSR. In Sec. 5.3, a probe tone much larger than an axion is used to characterize the squeezing over the bandwidth of the SSR. Section 5.4 describes a complete mock-haloscope experiment, injecting and detecting an axion-like tone of unknown frequency with and without the assistance of squeezing. Section 5.5 discusses protocols and calibrations common to the two experiments in this chapter. Section 5.6 then discusses complementary measurements to those taken in Sec. 5.3, and Sec. 5.7 describes the data processing used in Sec. 5.4. The chapter concludes in Sec. 5.8. Discussion then shifts from quantum to classical information processing, before ultimately returning to a real axion search, in which both tools are used advantageously.

Chapter 6 presents data processing and statistical analysis techniques commonly used in axion haloscopes. Section 6.1 overviews the processing protocol of the Haloscope At Yale Sensitive To Axion Cold dark matter (HAYSTAC) experiment, and Sec. 6.2 introduces some considerations pertinent to axion search analyses. Sections 6.3 and 6.4 describe frequentist frameworks commonly used to test for axions. In relation to these frameworks and that introduced in the next chapter, Sec. 6.5 discusses the role of the look-elsewhere effect in axion searches.

In Ch. 7, I introduce the Bayesian power-measured (BPM) statistical framework, a novel analysis designed for the particulars of axion direct detection, and demonstrate its utility on existing HAYSTAC data. Section 7.1 introduces the framework and describes key features pertinent to axion detection. Section 7.2 then evaluates the existing haloscope data processing protocols, and Sec. 7.3 contextualizes the BPM framework alongside the frequentist framework of Sec. 6.3, showing in what limits the two may be compared. Section 7.4 then applies the framework to the HAYSTAC Phase I dataset, achieving deeper reach into axion parameter space at zero hardware or operation cost.

The discussion of statistical frameworks for axion detection concludes in Sec. 7.5.

Chapter 8 combines the innovations of Chs. 5 and 7 to deliver a threefold combined speedup to the scan rate of Phase II of the HAYSTAC experiment. Section 8.1 describes the setup and basic operation of the experiment, Sec. 8.2 presents the data processing and analysis of the Phase II dataset, and Sec. 8.3 discusses squeezed state receiver operation and presents the data processing and results. We observe no signature of the axion, reporting exclusion results obtained at roughly three times the scan rate achievable without the innovations presented in this thesis. Section 8.4 discusses calibration measurements and protocols, and an issue with improper tuning rod thermalization during the data run is explored in Appx. C.

Lastly, Ch. 9 uses the BPM framework of Ch. 7 for a separate application, searching for the signature of axion-like particles (ALPs) in data from the JILA electron electric dipole moment (eEDM) experiment as much as 17 orders of magnitude lower in mass than the axions targeted in Ch. 8. Section 9.1 discusses how the JILA eEDM experiment may be sensitive to ALPs. Section 9.2 then uses a least-squares spectral analysis to extract spectral content of the measurement. Section 9.3 formulates and refines the competing ALP and no-ALP hypotheses, which are tested subject to the measured data, with results reported in Sec. 9.4 and prospects for future searches discussed in Sec. 9.5.

Chapter 10 concludes in Sec. 10.1, and in Sec. 10.2 discuss future prospects for speeding up dark matter axion searches both by using quantum measurement and by taking full advantage of the analysis framework presented in this thesis.

Chapter 2

Dark matter, axions, and haloscopes

Arguments pointing to the presence of more gravitationally interacting mass out in space than is indicated by visible light date back to Lord Kelvin in the late 19th century [1]. The ensuing hundred-plus years produced ever more compelling data points in support of the dark matter hypothesis: that a substantial fraction of the universe’s matter density, presently estimated at 85% [2], does not strongly interact with electromagnetism (or else we would see it, and it would not be “dark”), and is not made of particles from within the Standard Model of particle physics. Theoretical predictions and experimental evidence have each played integral roles in the history of dark matter, motivating promising theories and ruling some of them out. A dark matter candidate is viable insofar as it is well motivated by known physical phenomena, and not strongly disfavored by existing experimental evidence. Perhaps the dark matter candidate that best fits this description is the axion particle predicted to resolve an outstanding mystery in quantum chromodynamics (QCD), the study of the strong nuclear force.

This chapter presents an overview of the history of our understanding of dark matter. Starting with a broad focus, I address some of the key pieces of evidence for dark matter. Zooming in, I touch upon some of the most historically popular candidates to constitute the dark matter, and in particular the QCD axion. This chapter ends with a discussion of axion haloscopes, detectors designed to observe dark matter axions in the galactic halo by first converting them to microwave photons, and the present limitations upon their performance.

2.1 Evidence of dark matter

The case for dark matter has been built over the past century from sources ranging from the behavior of bodies within galaxies, to the interaction of galaxies, to measurements of the cosmos. This section discusses a few prominent examples: observations of the velocity dispersion of galaxies within clusters and of the speed of orbiting bodies within galaxies, which each provided some of the most compelling early evidence of dark matter. Evidence from the collision of two galactic clusters has strengthened the case for dark matter over competing theories, and, to date, the most precise measurements of dark matter density come from the cosmic microwave background (CMB) radiation [3]. The broader history of dark matter is a rich and fascinating topic, and the interested reader is encouraged to consult Ref. [4] for a more comprehensive overview or Ref. [5] for a nuanced examination of some of the sociological forces that drove physicists to the problem.

2.1.1 Galactic velocity dispersion in the Coma Cluster

Fritz Zwicky had one of the most impactful early insights into the composition of our universe when in the 1930's he first applied the virial theorem to the motion of luminous bodies within the Coma Cluster of galaxies [6, 7]. The virial theorem,

$$2\langle T \rangle = -n\langle V \rangle, \quad (2.1)$$

relates the total time-averaged kinetic (T) and potential (V) energies of a system bound by a potential of the form $V \propto r^{-n}$. For the gravitational potential energy

$$V_G = -\frac{Gm_1m_2}{r} \quad (2.2)$$

between point masses m_1 and m_2 a distance r apart, $n = 1$. Zwicky estimated $\langle T \rangle$ based on the observed amount of luminous matter within the cluster. He found an order-of-magnitude discrepancy between the velocity dispersion of galaxies implied by Eq. (2.1) and the actual, observed values.

Zwicky’s results implied a total abundance of matter ~ 500 times greater than just its luminous component. This result turns out to be incorrect in magnitude, but still remarkable for its prescience.² It would be decades before dark matter became fully recognized as one of the central, unexplained riddles of modern physics.

2.1.2 Galactic rotation curves

Fritz Zwicky’s early estimates of the dark matter density of the Coma Cluster had relied upon the velocity dispersion of galaxies, pointedly eschewing the possibility of inferring the mass of individual galaxies from their rotation curves: how fast luminous matter orbits the galactic center as a function of radial distance. Zwicky wrote: “It is not possible to derive the masses of [galaxies] from observed rotations, without the use of additional information” [4, 7]. The additional information Zwicky was referring to included the internal damping of motion within galaxies from the interaction of massive bodies. A few years later, however, Subrahmanyan Chandrasekhar showed that these viscous interactions were effectively negligible [9], opening up another window into gauging dark matter density.

Galactic rotation curves provide a very straightforward test of whether the matter distribution of a galaxy is as it appears — that is, proportional to the luminous matter. A simple Newtonian calculation (see e.g. Example 5.2 of Ref. [10]) shows that for a typical galaxy where the galactic mass M appears to be clumped tightly at the center, the orbital velocity v_{orb} should rise initially within the galaxy, and then fall off outside of it with radial distance r as

$$v_{\text{orb}} = \sqrt{\frac{GM}{r}}. \quad (2.3)$$

The dashed line in Fig. 2.1a shows what this would look like for the Andromeda galaxy. Spectrographic measurements (solid line) such as those taken by Vera Rubin and Kent Ford proved inconsistent with this, revealing a much flatter curve [11]. Measurements obtained using the 21

² In particular, Zwicky was using a now-obsolete value [8] for the Hubble constant, describing the rate at which the universe is expanding, to calculate the velocities of galaxies at a distance. An updated value [2] returns a more reasonable answer, and one still in line with the Zwicky’s basic conclusion.

cm line of hydrogen [12], Fig. 2.1b (left column), provided a proxy for the amount of luminous matter out at larger radii, where there are not many stars. While the luminous matter was indeed clumped towards the center across several observed galaxies, the rotation curves (right column) stayed remarkably flat past the point where most of the matter was contained. The flat rotation curves were suggestive of a dark matter halo around galaxies, much broader in extent than the directly observable cores of luminous matter [13].

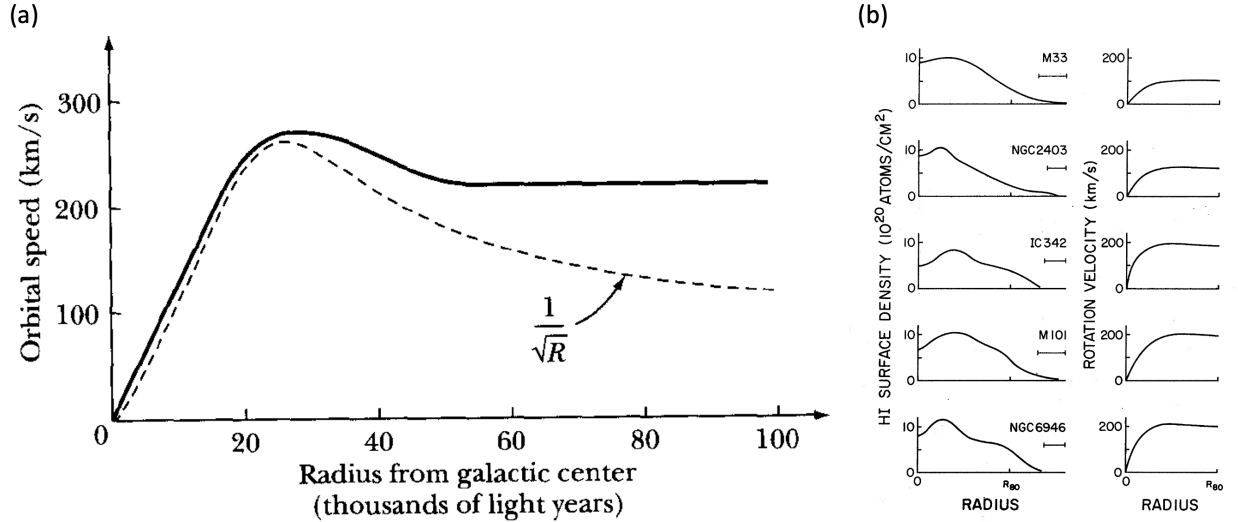


Figure 2.1: (a) Measurements of the orbital speed of the Andromeda galaxy (solid line) versus distance away from the galactic center shown against a simple model assuming only luminous matter, Eq. (2.3). (b) 21 cm line measurements (left column) of the hydrogen density within galaxies with flat rotation curves (right column). 80% of the hydrogen lies within R_{80} . Panel (a) reproduced from Ref. [10]. Panel (b) reproduced from Ref. [12].

2.1.3 Further arguments for and against dark matter

Arguments for dark matter have proliferated since the 1970’s, and have disfavored many otherwise plausible alternatives. For instance, it turns out that more matter exists as ionized hydrogen between galaxies than as stars within them [14]. Ionized hydrogen does not emit at the 21 cm line used in Sec. 2.1.2, instead emitting X-rays. This opens up the possibility that the “missing” matter revealed by galactic velocities could simply be made up of baryonic matter³ not

³ Baryonic matter is the name given to normal (i.e. non-dark) matter. It is understood to also include matter not composed of baryons (but still described by the Standard Model). While baryons are not the only type of “baryonic”

luminous at the frequencies being observed.

However, precise measurements of the anisotropy of the CMB weigh in heavily against this and other alternative explanations. The CMB radiation spectrum still contains the imprint of early universe physics, including the ratio between normal and dark matter. In essence, the galactic structure formation that we observe imprinted upon the anisotropy of the CMB could only have arisen from a very precise ratio of dark-to-normal matter, and is completely incompatible with the nonexistence of non-baryonic matter. These calculations come from a Λ CDM model accounting for normal matter, cold dark matter (CDM), and the cosmological constant (Λ) governing the universe's expansion [15].

One further popular class of alternatives to the dark matter hypothesis deserves remarking upon; these are theories of modified Newtonian dynamics (MOND), or modified gravity. First posited by Mordehai Milgrom in 1983 [16, 17], MOND posits that the gravitational force F_G may actually behave differently than the simple

$$F_G = \frac{Gm_1m_2}{r^2} \quad (2.4)$$

corresponding to Eq. (2.2). Of course, where local gravitational forces are strong, the modified force must reduce to Eq. (2.4). It is not hard to write down a functional form, such as

$$F_{G,\text{MOND}} = \frac{Gm_1m_2}{\mu\left(\frac{a}{a_0}\right)r^2} \quad (2.5)$$

with this property. Here, thinking of m_1 as the mass of a galaxy and m_2 as the mass of a star at some radius r from its center, μ is a function of m_2 's acceleration a relative to some new fundamental constant a_0 such that $\mu(a/a_0 \gg 1) \rightarrow 1$. Accounting for the flat rotation curves of Fig. 2.1 is as simple as positing that $\mu(a/a_0 \ll 1) \rightarrow a/a_0$.

One of several strong cases against MOND and for dark matter (in addition to the CMB) comes from observations taken in the early 2000's of a colliding pair of galaxy clusters [18, 19]. The Bullet Cluster is remarkable for being observed post-collision with a larger cluster of galaxies. An

matter, they do account for the vast majority of the baryonic matter density in the universe.

image of this collision is shown in Fig. 2.2. The figure highlights two regions of interest. First, X-ray emission from the galaxy clusters is shown in pink overlay. The X-rays are sourced from the large quantities of inter-galactic gases which account for the majority of the baryonic matter. The galaxy clusters also bend the path of light sourced from behind them, in a phenomenon known as gravitational lensing. The lensing effect reveals the location of the gravitationally attracting matter, shown in blue overlay. The key feature of Fig. 2.2 is that for both clusters the blue and pink regions are spatially separated. This is easy to account for with dark matter: the hot, ionized gas between galaxies is slowed as it interacts electromagnetically with that of the other cluster, while the dark matter sails through unimpeded. Though there remain a few voices of dissent [20], the weight of evidence from the Bullet Cluster seems to favor the dark matter explanation over theories of modified gravity [21].



Figure 2.2: Optical image of the Bullet Cluster collision with overlaid X-ray and lensing map. Sources of X-ray emission within the clusters indicate the location of baryonic matter (pink). The collision slowed down the intra-cluster gasses, spatially separating them from the dark matter region (blue), which has been reconstructed from gravitational lensing studies. The separation of the luminous matter from the center of gravity is strong evidence for dark matter. Figure reproduced from Ref. [22].

The evidence discussed above represents just a small sampling of some of the most historically

important advances in our understanding of the dark matter problem, and are by no means a comprehensive survey thereof. Moving forward, we will ask what dark could be composed of, and what its essential properties are.

2.2 The not-so-miraculous WIMP

The fields of particle physics and astrophysics motivate (or in some cases merely allow) particles or objects that interact only weakly (or strongly, but from our perspective occasionally) with normal matter, and could plausibly populate much of space. Some of the most popular early candidates such as massive astrophysical compact halo objects (MACHOs) have been ruled out by observational studies [23, 24]. Others, including weakly⁴ interacting massive particles (WIMPs) [25] have had meaningful fractions of formerly favored parameter space excluded by laboratory experiments.

WIMPs are the lightest stable particles predicted by supersymmetry (SUSY) theories, and are well motivated from both theoretical and experimental perspectives [26]. Their popularity traces to two sources. First is the “WIMP miracle,” the fact that the number of WIMPs expected to have survived until the present day should decrease with their self-annihilation cross-section $\langle\sigma v\rangle$, predicted to be of order $3 \times 10^{-26} \text{ cm}^3/\text{s}$ to account the present dark matter density⁵ — and this cross section is the right scale for a weakly interacting particle with mass $m_\chi \sim 100 \text{ GeV}/c^2$. Second, we know how to detect them. WIMPs are predicted to scatter off of baryonic matter, producing sound [28], light, and/or charge signals in physical detectors, which often seek to correlate events in two of these three categories to veto backgrounds.

These experiments have been successful in sweeping out much of the WIMP parameter space shown in Fig. 2.3 [30–32]. The WIMP-baryon scattering cross-sections originally predicted by the WIMP miracle have been ruled out, weakening the case for WIMPs. Additionally, as detectors become more and more sensitive, they approach the “neutrino floor” (orange shaded region in

⁴ “Weakly” in this acronym refers to the weak nuclear force, and not the feeble coupling to normal matter, though neither meaning would be a mischaracterization.

⁵ So long as $m_\chi \gtrsim 10 \text{ GeV}/c^2$, the predicted value for $\langle\sigma v\rangle$ is not strongly dependent on m_χ [27].




Figure 2.3: Weakly interacting massive particle (WIMP) parameter space showing spin-independent (SI) WIMP-nucleon cross section vs. WIMP mass. Experimental exclusion results are rapidly closing in on the neutrino floor (orange shaded region), where coherent scattering of neutrinos will make it difficult to detect a WIMP dark matter signal. See Ref. [29] for references to experimental results. Figure reproduced from Ref. [29].

Fig. 2.3), where ubiquitous neutrino scattering events will provide a formidable background [33]. WIMP direct detection in the presence of neutrino-scattering may still be achievable [34], but will prove more difficult. The theoretical and experimental cases for WIMPs, while by no means inviable [35], are thus each less compelling than they seemed in the 1970's and 1980's, and, as such, the case for competing hypotheses is strengthened.

2.3 The QCD axion solution to dark matter

A particularly compelling candidate for dark matter is the QCD axion. We will see that the QCD axion, like the WIMP, solves an outstanding mystery of modern physics and possesses key attributes of dark matter. Unlike the WIMP, experimental evidence has not yet weighed in strongly on the most plausible parameter space.

2.3.1 The strong CP problem

The QCD axion was “invented” to resolve a mystery seemingly unrelated to that of dark matter: why the strong nuclear force seems to obey charge-parity (CP) symmetry [36, 37].⁶ The problem was vexing because CP symmetry is known not to be a good symmetry of nature. CP violation was first observed in neutral kaon decays [39], and it took almost four decades before it was observed outside of kaon physics, in B mesons [40]. Recently, long-baseline measurements of neutrino oscillations observed CP violation in elementary particles for the first time [41]. But what all of these occurrences have in common is that they are governed by the electroweak interaction.

In the strong sector, CP violation has never been observed, and not for a lack of sensitivity to it. CP violation in the strong interaction would be directly observable via the measurement of a nonzero neutron electric dipole moment (nEDM) d_n . In particular, d_n scales as a parameter $\bar{\theta}$ arising in part from QCD.⁷ However, the most precise nEDM measurement to date of $d_n = (0.0 \pm 1.1) \times 10^{-26} \text{ ecm}$ [43] accommodates $\bar{\theta} < 10^{-10}$, whereas its being an otherwise unconstrained angle suggests it should be $\mathcal{O}(1)$.

This part-in- 10^{10} fine-tuning problem has several physically allowed explanations, the most plausible of which predicts QCD axion particles as a byproduct. First, it could truly just be random chance. But this is, by construction, unlikely, and there is no anthropic necessity for $\bar{\theta}$ to be zero [42]. Second, CP symmetry could actually be an exact symmetry of nature that is spontaneously broken [42]. But explanations of this sort tend to either be ruled out by existing measurement [29] or require fine-tuning of their own atop elaborate theoretical mechanisms [44–46]. Third, one of the quarks being massless would make d_n independent of $\bar{\theta}$, but current measurements of the up quark (the lightest quark) place its bare mass at $2.2_{-0.4}^{+0.5} \text{ MeV}/c^2$ [29], inconsistent with

⁶ CP symmetry exists if, upon flipping the signs of all the charges in the universe and inverting the positions of all particles in space about some point, a careful observer cannot tell the difference between the flipped universe and the original. Charge-parity-time (CPT) symmetry, which is symmetry under a CP operation plus a reversal of the arrow of time, is believed to be a good symmetry of nature [38]. Therefore, CP violation implies T violation: the universe does not obey the same laws of physics with the tape run backwards.

⁷ Another contribution to $\bar{\theta}$ actually comes from a Yukawa term in the electroweak interaction, making the so-called “strong” CP problem a problem for the whole Standard Model: $\bar{\theta} = \theta_{\text{QCD}} + \theta_{\text{Yukawa}}$. For a far more detailed explanation of the attendant physics, see Ref. [42].

zero. Finally, there is the Peccei-Quinn (PQ) solution to the strong CP problem.

2.3.2 The Peccei-Quinn mechanism

In 1977, Roberto Peccei and Helen Quinn proposed a theory in which $\bar{\theta}$ is promoted from a static parameter to a dynamical field, causing it to naturally relax to 0 and elegantly cleaning up the strong CP problem [36,37].⁸ This cleanliness-promoting feature gave rise to the name “axion,” after the dish detergent [48]. Axions are pseudo-Goldstone bosons⁹ that arise as an immediate consequence of the PQ mechanism [49,50]. It was not long before particle physicists realized that sufficiently light axions have key features resembling dark matter. This main results of this thesis will pertain to axions of mass $m_a \sim 10 \mu\text{eV}/c^2$.

2.3.3 Axions as dark matter candidates

The light axions we will focus on are bosonic, stable on cosmological timescales, and weakly interacting. They also may have a production mechanism capable of making up in numbers what the individual particles lack in mass, while allowing for them to be nonrelativistic. We will see that all of these features together make axions a promising dark matter candidate.

Any sufficiently light dark matter candidate must be bosonic. This is because fitting enough fermions into a galaxy to account for observed dark matter densities gives the highest energy particles a Fermi velocity greater than galactic escape velocities. Their low masses [29] are one way we know that neutrinos, for example, do not constitute dark matter [42]. Axions are indeed bosons, meeting this criterion.

To be viable, a dark matter candidate must also have been cold — that is, nonrelativistic — at the time of galactic structure formation. In particular, baryonic matter only begins to form large-scale structure after the decoupling of radiation from protons and electrons, which occurs

⁸ For a delightful analogy illustrating the PQ solution to the strong CP problem, see Ref. [47].

⁹ While Goldstone bosons are the massless excitations of a flat potential resembling the circular mode of travel at the bottom of a wine-bottle, axions are *pseudo*-Goldstone bosons because their wine bottle-potential is tilted. The mass of a particle scales with the square root of the curvature of this potential, so axions have nonzero, but still relatively low, masses.

when the universe is $\sim 380,000$ years old. In a universe with only baryonic matter, this would not allow enough time for the structure we observe to have formed [42]. Dark matter, meanwhile, can in principle begin clumping once the universe become matter dominated, at $\sim 50,000$ years. However, lighter particles usually have a higher velocity. The equipartition theorem gives the root-mean-square (RMS) velocity of particles of mass m in an ideal gas at temperature T as

$$v_{\text{RMS}} = \sqrt{\frac{3k_B T}{m}}, \quad (2.6)$$

where k_B is the Boltzmann constant. The faster velocity gives rise to a longer free-streaming length, the characteristic distance that the particles travel before being slowed by cosmic expansion. Structure in matter density fluctuations in the early universe at smaller spatial scales than this free streaming length is wiped out. Dark matter with primordial RMS velocities comparable to the speed of light, $v_{\text{RMS}} \sim c$, precludes the structure needed to form galaxies through the local clumping of matter beginning at $\sim 50,000$ years. “Hot,” or relativistic, dark matter has accordingly been shown through simulations to form galaxies very slowly, producing a universe different than the one we find ourselves in [51]. If the dark matter were baryonic, it would be in equilibrium with radiation and thus hot.

Axions, despite likely being even lighter than neutrinos, can be cold if they were produced athermally, that is: not in thermal equilibrium with their early-universe surroundings. The T in Eq. (2.6) is no longer required to be the universe’s temperature at the time of structure formation, and light dark matter is allowed provided it is sufficiently abundant. In 1983, several independent groups of theorists [52–54] showed how PQ mechanism may account for dark matter by athermally producing axions. Furthermore, this mechanism produces *disproportionately* more axions the lighter m_a is.¹⁰

We have seen that axions could have been produced in large quantities as cold bosons in

¹⁰ A key determinant of the number of axions produced derives from whether a spontaneous symmetry breaking at the heart of the PQ mechanism occurred before or after the early universe underwent inflation. If symmetry breaking occurs before inflation, the quantity of axions produced is in a sense much more random: it scales as the absolute value of random angle going from $-\pi$ to π . If symmetry breaking occurs afterwards, many independently chosen values are averaged over, eliminating that source of uncertainty. Sufficiently low-mass QCD axions ($m_a \lesssim 10\mu\text{eV}/c^2$) require symmetry breaking before inflation, plus some fine-tuning for the angle.

the early universe; what remains to qualify them as a dark matter candidate is their interaction with normal matter. Axions are predicted to interact weakly with most forms of normal matter. Though they have no electromagnetic charge, their dominant interaction is with electromagnetism via a two-photon coupling in the universe's Lagrangian,

$$\mathcal{L} \supset g_{a\gamma\gamma} |\mathbf{E} \cdot \mathbf{B}| a, \quad (2.7)$$

where \mathbf{E} and \mathbf{B} are the local electric and magnetic field vectors, respectively, and a is the axion field amplitude. We will henceforth reparametrize the axion-photon coupling $g_{a\gamma\gamma}$ as an order-unity dimensionless number

$$g_\gamma = \frac{g_{a\gamma\gamma}}{m_a c^2} \left(\frac{\pi \Lambda^2}{\alpha} \right), \quad (2.8)$$

where $\alpha \approx 1/137$ is the fine structure constant and $\Lambda = 77.6$ MeV encodes the dependence of the axion mass on hadronic physics. We have normalized out the linear scaling with rest energy $m_a c^2$, because QCD axion couplings to normal matter are proportionate to m_a .¹¹ The value that g_γ takes is model-dependent, with the benchmark Kim-Shifman-Vainshtein-Zakharov (KSVZ; $g_\gamma^{\text{KSVZ}} = -0.97$) [55, 56] and Dine-Fischler-Srednicki-Zhitnitsky (DFSZ; $g_\gamma^{\text{DFSZ}} = 0.36$) couplings [57, 58] serving as more and less optimistic targets, respectively, for detection experiments.¹² These values for g_γ imply an extremely long lifetime for the axion with respect to its two-photon decay channel: axions below the ~ 10 eV mass scale are stable on cosmological timescales [42]. Moreover, for any axion in the QCD band [59] containing both KSVZ and DFSZ models, g_γ is sufficiently small that our nonobservance of axions thus far is not meaningful evidence against their existence. Rather, if QCD axions — which have all the essential properties of dark matter — in fact constitute dark matter, it will take a concerted effort to find them.

¹¹ More precisely, $g_{a\gamma\gamma}$ and m_a both scale inversely as the energy scale f_a at which PQ symmetry-breaking occurs. This energy scale dwarfs those achieved by physicists in high energy experiments, including the 13 TeV record at the Large Hadron Collider. The largeness of f_a is why the axion is predicted to be both so light and so weakly coupled.

¹² Physically, g_γ can take positive or negative values, but since only g_γ^2 contributes to the axion signal power, I will henceforth use g_γ to stand in for $|g_\gamma|$ throughout this thesis.

2.4 Detectable features of axionic dark matter

The properties discussed above qualify axions as dark matter candidates; now I will discuss properties pertinent to their detection. A simple yet well motivated and actionable view of axions holds that they exist in a bounded, “1.1-dimensional” parameter space of mass m_a and coupling g_γ with known velocity distribution and density throughout the cosmos. To start, let us consider the axion-photon parameter space in which this thesis will primarily work.

2.4.1 Axion parameter space

Figure 2.4 shows a portion of axion-photon parameter space which has attracted attention in recent years for its experimental accessibility [60]. Most of the plot (white and yellow regions) has not been experimentally probed. The diagonal yellow region represents the QCD axion band,¹³ with the KSVZ line having unit slope in the space of $g_{a\gamma\gamma}$ and m_a . While non-QCD axions, also known as axion-like (ALPs),¹⁴ of very high couplings would have imprinted a signal in devices such as helioscopes (excluded blue region) [61], the only demonstrated method for reaching QCD couplings is with axion haloscopes (excluded green regions), discussed below.

The parameter space between the dashed, black lines at $10 \mu\text{eV}/c^2$ and $1 \text{meV}/c^2$ in Fig. 2.4 is theoretically favored by a cosmological and an astrophysical argument. Axions below the $m_{\text{OC}} \sim 10 \mu\text{eV}/c^2$ scale in models in which PQ symmetry-breaking occurs after inflation tend to be overproduced to the point of “overclosing” the universe: that is, more dark matter is predicted than is observed, disfavoring lower masses [52]. In the event that inflation did occur first, QCD axions are allowed both below and above m_{OC} , making it a somewhat soft bound.

The upper bound of $m_{\text{SN}} \sim 1 \text{meV}/c^2$ is set by the $\sim 10 \text{s}$ temporal duration of the observed neutrino burst of supernova 1987A [63]. The production of axions in the exploding star should scale with their coupling to normal matter, which in turn scales as m_a . Sufficiently heavy axions

¹³ Different physicists draw the boundary of the QCD band slightly differently, though everyone agrees it encompasses both the KSVZ and DFSZ lines.

¹⁴ ALPs, which will feature in Ch. 9, do not solve the strong CP problem. Since there are no strong bounds on the ALP-photon coupling, it is also possible that ALPs could lie below the QCD band, which would make them very difficult to detect.




Figure 2.4: A portion of dark matter axion-photon parameter space prior to the work discussed in this thesis. The white and diagonal yellow regions are unprobed, with the diagonal yellow region highlighting the theoretically favored QCD band. Haloscope exclusion is shown in green, with other excluded regions shaded different colors. For a thorough discussion of the existing constraints within axion parameter space, see Ref. [60]. The black, dashed lines show the $10 \mu\text{eV}/c^2$ – $1 \text{meV}/c^2$ range of axion masses favored by overclosure arguments and supernova 1987A bounds discussed in the text. The orange, dashed lines enclose the 1–10 GHz range of axion mass-converted frequencies ν_a , Eq. (2.9), of particular interest to this thesis. See Ref. [29] for references to experimental results. Figure adapted from Ref. [62].

would have made off with much of the star’s energy, leaving fewer neutrinos to lengthen the burst duration by colliding with each other. This is considered by many a more concrete bound than that of overclosure, but it is conceivable that modelers have misunderstood the origin of the observed neutrinos, which would undo the argument [64]. But even if one is equally drawn to the adjacent regions, the two decade window of roughly $10 \mu\text{eV}/c^2$ – $1 \text{meV}/c^2$ remains well-motivated parameter space.

Zooming farther into the $10 \mu\text{eV}/c^2$ – $1 \text{meV}/c^2$ window, arguments abound for specific favored or disfavored mass ranges [65–68]. Of particular relevance to the work presented in this thesis, Klaer and Moore ($26.2 \pm 3.4 \mu\text{eV}/c^2$) [69] and Buschmann, Foster, and Safdi ($25.2 \pm 11.0 \mu\text{eV}/c^2$) [70] have independently predicted the axion to lie in overlapping regions around $25 \mu\text{eV}/c^2$. Near and

below these masses, the axion acts more like an oscillating field at with axion rest-mass converted frequency (henceforth “axion frequency”)

$$\nu_a = \frac{m_a c^2}{h}. \quad (2.9)$$

It is beyond the scope of this work to adjudicate between theoretical models, except by contributing to the experimental toolkit and the body of axion direct detection data that will enable us to answer the question directly.

2.4.2 Local axion dark matter density

In order to discuss how axions may be detected, two final features of cold dark matter bear examining; the first is the local dark matter density ρ_{DM} . In this work and throughout the axion exclusion literature, experimental constraints such as those in Fig. 2.4 are typically placed on the quantity g_γ , or equivalently $g_{a\gamma\gamma}$ via Eq. (2.8). Implicit in these constraints are two assumptions: that axions saturate 100% of the local dark matter density ($\rho_a = \rho_{\text{DM}}$) and that the local dark matter density takes a specific value, typically quoted as $\rho_{\text{DM}} = 0.45 \text{ GeV/cm}^3$. The observable combination that shows up in experiments is generally $\rho_a g_\gamma^2$, so it is a trivial matter to scale any constraint placed on g_γ to assume a different value of ρ_a . For consistency with the literature, in this thesis we will use $\rho_a = 0.45 \text{ GeV/cm}^3$ unless otherwise stated.¹⁵

2.4.3 Axion lineshape

The other feature of axionic dark matter that will prove relevant to axion detection is the total energy (E_{tot}), or equivalently frequency $\nu = E_{\text{tot}}/h$, of the axion particles.¹⁶ Here, we will not want to think of ν as a constant, but rather a random variable: a quantity that can take a range of values specified by a probability density function $f_\nu(\nu)$, or equivalently its integral, the

¹⁵ WIMP dark matter experiments have used the more conservative value $\rho_{\text{DM}} = 0.3 \text{ GeV/cm}^3$, also consistent with measurement [71]. Though the primary goal is to detect dark matter, the secondary goal of certifying that it does not exist in a specified coupling range does seem more naturally served by the more conservative value. In Ch. 9, for consistency with the ALP literature, we will use 0.40 GeV/cm^3 .

¹⁶ Thinking of axions as particles will be helpful here. The results will work just as well for measurements that take advantage of the wave-like properties of axions.

cumulative distribution function $F_\nu(\nu) = \int_{-\infty}^{\nu} f_\nu(\nu') d\nu'$. The frequency distribution is obtained through approximating the total energy as the rest energy plus the nonrelativistic kinetic correction:

$$E_{\text{tot}} \approx m_a c^2 + \frac{1}{2} m_a v^2. \quad (2.10)$$

For cold dark matter, the dimensionless velocity $\beta = v/c$ is $\mathcal{O}(10^{-3})$, meaning the nonrelativistic term is of order 10^6 times greater than the first relativistic correction $3m_a v^4/8c^2$ to Eq. (2.10), making this an excellent approximation. Determining the frequency distribution then amounts to determining the axion frequency ν_a and the velocity distribution $f_v(v)$ of axion particles. The axion frequency is a free parameter of our search, and assumed to have a single, constant value, hence $f_\nu(\nu) \rightarrow f_\nu(\nu|\nu_a)$. The velocity distribution is model-dependent, and in this thesis we will stick to the standard halo model (SHM), which describes virialized axions as observed from the moving laboratory frame of our solar system.¹⁷

The virial theorem, Eq. (2.1), relates the mean kinetic energy $\langle T \rangle = m_a \langle v^2 \rangle / 2$ and gravitational potential energy $\langle V_G \rangle$ of axionic dark matter in the Milky Way Galaxy, allowing us to write down $f_v(v)$ and ultimately $f_\nu(\nu)$ in the galactic rest frame. The galactic rest frame is defined as the frame in which our galaxy has no net angular momentum. In it, the velocity components $\{v_x, v_y, v_z\}$ of axions are independent, identically distributed Gaussian random variables with mean zero and variance $\sigma_{v_x}^2 = \langle v^2 \rangle / 3$: $\{v_x, v_y, v_z\} \sim \mathcal{N}(0, \sigma_{v_x}^2)$. Summing the velocity component random variables in quadrature yields the Maxwell-Boltzmann (MB, or Maxwellian) distribution for the velocity amplitude of particles in an ideal gas:

$$f_{v,\text{MB}}(v) = \frac{4}{\sqrt{\pi}} \frac{v^2}{v_c^3} \exp \left[- \left(\frac{v}{v_c} \right)^2 \right], \quad (2.11)$$

where $v_c = \sigma_{v_x} / \sqrt{2} \approx 220$ km/s is the modal velocity of the distribution. The corresponding MB frequency distribution obtained using Eq. (2.10) is

$$f_{\nu,\text{MB}}(\nu|\nu_a) = 4 \sqrt{\frac{2}{\pi}} \left(\frac{c}{v_c} \right)^3 \frac{1}{\nu_a} \sqrt{\frac{\nu - \nu_a}{\nu_a}} \exp \left[-2 \left(\frac{c^2}{v_c^2} \right) \left(\frac{\nu - \nu_a}{\nu_a} \right) \right] \quad (2.12)$$

¹⁷ The techniques discussed in this work are applicable independent of axion lineshape, and may be readily adapted as new evidence refines our knowledge of the dark matter halo [72, 73], or to test alternative lineshapes [74, 75].

for $\nu_a \geq \nu$ and 0 otherwise. The axion frequencies would be MB-distributed according to Eq. (2.12) if our solar system were stationary with respect to the galactic rest frame. Instead, we move at $v_s \approx 232$ km/s through a sea of axions in the galactic rest frame. The resulting SHM velocity and frequency distributions, results which will be of direct use to us in Chs. 8 and 9, are

$$f_{\nu,\text{SHM}}(v) = \frac{v}{\sqrt{\pi}v_c v_s} \exp \left[- \left(\frac{v + v_s}{v_c} \right)^2 \right] \left(\exp \left[\frac{4vv_s}{v_c^2} \right] - 1 \right) \quad (2.13)$$

and

$$f_{\nu,\text{SHM}}(\nu|\nu_a) = \frac{2}{\sqrt{\pi}} \left(\frac{c^2}{v_c v_s} \right) \frac{1}{\nu_a} \sinh \left(\frac{v_s c}{v_c^2} \sqrt{8 \frac{\nu - \nu_a}{\nu_a}} \right) \exp \left[-2 \left(\frac{c}{v_c} \right)^2 \frac{\nu - \nu_a}{\nu_a} - \left(\frac{v_s}{v_c} \right)^2 \right], \quad (2.14)$$

again only for $\nu \geq \nu_a$ [75, 76]. The lineshapes, Eqs. (2.12) and (2.14), are plotted in Fig. 2.5. The primary effect of $v_s > 0$ is to broaden¹⁸ $f_{\nu,\text{SHM}}$ (blue) relative to $f_{\nu,\text{MB}}$ (green), but in either case the fractional linewidth $\Delta_a \sim \nu_a 10^{-6}$ (equivalently, the quality factor of the axion is of order a million: $Q_a \sim 10^6$).¹⁹

This concludes our brief overview of axion theory. For a far more thorough discussion of much of what was covered here, the interested reader is encouraged to consult Ref. [42]. We are now ready to discuss how it is that QCD axions may be detected.

2.5 The axion haloscope

The axion haloscope was first proposed by Pierre Sikivie in the 1980's to detect dark matter axions in the galactic halo by converting them to photons²⁰ in a microwave cavity using the inverse Primakoff effect and a large static magnetic field $\mathbf{B} = B_0 \hat{\mathbf{z}}$ [85–87]. The cavity delivers a Q -factor enhancement so long as the electric field \mathbf{E} of at least one of its modes overlaps with $\hat{\mathbf{z}}$ so as to align with \mathbf{B} , Eq. (2.7). The prototypical haloscope uses a right circular cylindrical cavity with

¹⁸ Narrower lineshapes such as those given by Eq. (2.12) or Ref. [74] are in principle easier to detect because of their higher signal-to-noise ratios [77].

¹⁹ Further corrections come from the orbital and rotation motion of the earth around the sun and its axis, respectively. For virialized axions these corrections are negligible [76, 78].

²⁰ What is more precisely meant by “photons” unless otherwise indicated is “electromagnetic energy.” While the energy could in principle be measured in the number basis [79–82], at the frequencies where this thesis focuses, it is more efficient to measure the continuous quadrature observables [83, 84]. It is nonetheless common to refer to axion-induced excesses in a haloscope cavity as photons.




Figure 2.5: Virialized axion lineshapes $f_\nu(\nu|\nu_a)$ for $\nu_a = 4$ GHz. The Maxwell-Boltzmann (MB) spectral distribution, Eq. (2.12) (green), is narrower than the standard halo model (SHM) distribution, Eq. (2.14) (blue), which accounts for the motion of the sun in the galactic rest frame. The means μ and standard deviations σ of the two distributions are indicated by dots and bars, respectively.

mode indices m , n , and l ²¹ for which the spatial overlap associated with a given cavity mode is parametrized as the form factor

$$C_{mnl} = \frac{(\int_V \mathbf{E}(\mathbf{x}) \cdot \hat{\mathbf{z}} d^3\mathbf{x})^2}{\int_V \varepsilon(\mathbf{x}) |\mathbf{E}(\mathbf{x})|^2 d^3\mathbf{x}}, \quad (2.15)$$

where V is the volume of the cavity and $\varepsilon(\mathbf{x})$ is the position (\mathbf{x})-dependent permittivity.²² Figure 2.6a shows the field directions (ignoring amplitude) of transverse magnetic (TM), transverse electric (TE) and transverse electromagnetic (TEM) modes for simple cylindrical cavities. For a cylindrical cavity, the only modes that produce a nonvanishing form factor are the TM_{0n0} modes [88, 92]. However, an unperturbed cylindrical cavity is of limited use for axion detection because its TM_{010} mode — the one with the highest form factor and lowest density of “intruder” modes insensitive to

²¹ These correspond respectively to the m , n , and p indices in Sec. 8.7 of Ref. [88].

²² For the haloscope central to this thesis, $\varepsilon(\mathbf{x}) = \varepsilon_0$, the vacuum permittivity, everywhere. Some haloscope experiments have taken advantage of the dielectric properties of liquid helium [89] or dielectric rods [90, 91] to tune their cavity frequencies.

the axion — can only probe around a fixed resonant frequency,²³ and broadening its resonance to achieve a larger bandwidth incurs disproportionate losses in sensitivity [95]. The usual solution is to insert a metal and/or dielectric tuning rod coaligned with the vertical axis $\hat{\mathbf{z}}$ of the cavity. For a metal rod, rotation away from the center of the cavity, shown in Fig. 2.6b, tunes the TM_{0n0} -like modes to lower frequencies.²⁴

2.5.1 Previous haloscope results

Axion haloscopes have been successfully running since the late 1987, but have as yet seen no evidence of the axion. Figure 2.7 shows in greater detail than Fig. 2.4 the region of parameter space that has been partially excluded by haloscopes.²⁵ The earliest experiments were carried out at Brookhaven National Laboratory in collaboration with the University of Rochester and Fermilab [97], and at the University of Florida [98]. The Axion Dark Matter Experiment (ADMX) came online in 1998 [89], and between then and 2016 [99] was the only operational haloscope, primarily probing parameter space at KSVZ sensitivity.

In 2017, the Haloscope At Yale Sensitive To Axion CDM (HAYSTAC) became the second active haloscope [90, 106], sensitive to higher couplings g_γ and operating at higher frequencies ν_a than ADMX. We will see in Sec. 2.5.2 that going deep in coupling and high in frequencies are comparably difficult propositions. As of early 2020, the division of labor in the axion detection field featured one experiment actively pursuing each of these objectives, with ADMX sensitive to \sim DFSZ axions around 660 MHz [91], and HAYSTAC searching for \sim 3-KSVZ axions around 5.7

²³ While most axion haloscopes have only used the TM_{010} mode, the ADMX Sidecar and ORGAN experiments recently demonstrated use of the TM_{020} mode to collect axion-sensitive data [93, 94]. Using the TM_{020} mode (or even higher modes) affords sensitivity to higher frequencies with a much higher cavity volume V than would be achievable with the TM_{010} mode of a smaller cavity. However, the tradeoff is not as favorable as it may appear: the electric field undergoes 1 1/2 oscillations over the height of the cylinder, and the form factor integral, Eq. (2.15), ensures that the positive and negative half of the first oscillation cancel out, reducing the form factor roughly threefold as the volume grows by the same amount. Furthermore, around higher TM_{0n0} mode frequencies there is a higher number of unwanted crossings with axion-insensitive cavity modes. The real utility of higher modes is that in principle they can be read out simultaneously, opening up parallel data streams at no cost to the first.

²⁴ The metal rod has the side effect of reducing the available cavity volume used for axion conversion. The effect is not always a small one: in HAYSTAC, the rod occupies roughly 1/4 of the cavity volume [96], reducing the experiment's main figure of merit, its scan-rate, roughly twofold.

²⁵ Excluding with a haloscope at a given frequency and coupling entails excluding all stronger couplings at the same frequency.

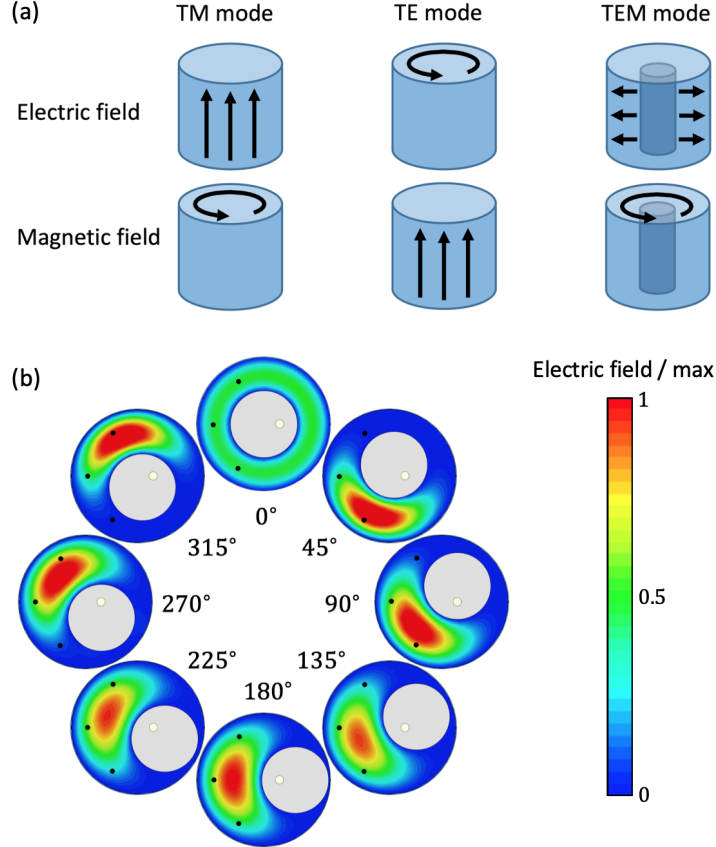


Figure 2.6: (a) Electric and magnetic field directions for simple cylindrical cavity geometries. Only the subset of transverse magnetic (TM) modes whose electric fields integrate to a nonzero value along the vertical direction (aligned with the magnetic field) are useful for axion detection. (b) Simulated electric field magnitudes for the TM_{010} -like mode of a cylindrical cavity with a metal tuning rod at different angles. The changing mode structure tunes the resonance frequency and alters the form factor C_{010} , Eq. (2.15). Figure adapted from Ref. [62].

GHz [107].

Additionally, the CAPP-8TB experiment has reported its first results at 1.6 GHz [108], the ADMX Sidecar experiment has probed parameter space above the QCD band at a range of higher frequencies [93], and the QUAX- $a\gamma$ [109] and ORGAN [94] experiments have reported early results at much higher frequencies, but with limited sensitivity. Concurrent with these developments, other haloscope-like platforms have been [110–115] coming online or are poised to do so [116, 117].²⁶

²⁶ These experiments all look for axions on the assumption that they constitute dark matter. However, axions could very well exist and not constitute any meaningful portion of dark matter. In that event, “light shining through walls” (LSW) experiments that attempt to first create and then detect axions [118–120] (upper yellow region in Fig. 2.4), axion helioscopes which look for axions produced in the sun [61, 121, 122] (blue region in Fig. 2.4), or searches for




Figure 2.7: Axion-photon exclusion obtained using haloscopes prior to the main work of this thesis. The regions labeled RBF (pink) and UF (blue) represent early results insensitive to the QCD axion (yellow band) obtained at Brookhaven National Laboratory [97, 100] and University of Florida [98, 101], respectively. ADMX (light red) [89, 91, 93, 99, 102–105], HAYSTAC (green) [106, 107], and CAPP (red) [108] have all reached the QCD band.

The collective null result of Fig. 2.7 only extends over a minute fraction of the parameter space axions could plausibly inhabit. Nevertheless, there is real cause for optimism: haloscope searches have sped up roughly one millionfold between the first search carried in 1987 [97] and the major results of the past few years [91, 105–107]. Such comparisons are made between the scan rates of axion haloscopes.

2.5.2 Haloscope scan rate

Scan rate R is the primary figure merit for axion haloscopes:²⁷ it quantifies how rapidly a haloscope can scan through parameter space at some fixed g_γ . It has the curious dimensions of

axion-mediated spin forces [123, 124] may be the best hopes to resolve the strong CP problem.

²⁷ This is admittedly a somewhat subjective assessment, as it is hard to compare the more objective scan rate with other important haloscope traits, such as affordability, ease of operation, density of intruder modes, and accessible parameter space. Furthermore, because haloscopes spend nonnegligible amounts of time rescanning potential signals at individual frequencies, the maximum single-frequency sensitivity, which can be different for two haloscopes of the same scan rate, is an important secondary figure of merit. For the purpose of this thesis, I will treat improving scan rate as the ultimate goal, paying attention where appropriate to other important objectives. This simple framing is very useful, and can be made to encompass some of the others figures of merit mentioned above.

frequency per time, and is proportionate to several key parameters [42]:

$$R \propto \left(\frac{g_\gamma^4}{\text{SNR}'^2} \right) \left(\frac{\rho_a^2 Q_a}{\Lambda^8} \right) \left(\frac{B_0^4 V^2 C_{mnl}^2 Q_0}{N_{\text{sys}}^2} \right), \quad (2.16)$$

The first set of parentheses consists of scan properties chosen by the experimentalist: what signal-to-noise power ratio (SNR) is desired at what coupling g_γ for the experiment. The prime on the SNR is to denote a modification that owes to the look-elsewhere effect, and is derived in Sec. 6.5. We will see that as a rough approximation, the denominator of the first term can be thought of as being

$$\text{SNR}'^2 \sim \text{SNR}^2 + 2 \log(N_I), \quad (2.17)$$

where N_I is the number of “independent” frequencies being probed for the axion. The second set of parentheses comprises several physical parameters of axion dark matter. The final set contains experimental parameters: the magnetic field B_0 , cavity volume V , form factor C_{mnl} , and unloaded cavity quality factor Q_0 ²⁸ should be maximized,²⁹ while the total system noise N_{sys} (in units of photons, or quanta) should be minimized. Improving a haloscope’s scan rate through hardware is typically synonymous with improving one of these parameters.³⁰

2.5.3 Improving the scan rate

The millionfold improvement in scan rate accrued over the last three decades owes predominantly to dramatic advances in just one of the experimental parameters in the final parenthetical of Eq. (2.16): namely, the total system noise N_{sys} . Cylindrical cavity volumes change as

$$V \propto \nu^{-\zeta} \quad (2.18)$$

²⁸ Equation (2.16) assumes Q_0 to be smaller than Q_a , as has been the case for all haloscopes to date. Improving Q_0 beyond Q_a while operating in the standard manner does not improve scan rate farther, as the coherence time of the axion becomes the limiting factor.

²⁹ We will see in Ch. 5 that another parameter of importance to the scan rate is how strongly the cavity is coupled to its readout chain. This parameter is in practice easy to adjust to its optimal value, so for now we leave it out of Eq. (2.16), which only includes quantities which pose fundamental or practical limitations to scan rate.

³⁰ Scan rate can be improved in other ways as well. Advances in analysis techniques [75, 78, 125, 126] that allow for more efficient use of information will be a central focus of this thesis. Additionally, improvements to acquisition live time can come through either improved software or hardware. Optimizing scan strategy remains a largely unexplored avenue by which total scan time might be reduced [95, 126].

to accommodate one half-wavelength at the frequency of interest, where $\zeta \in [2, 3]$ can be as low as 2 if experimentalists are able to tolerate the linear growth in TE mode density with cavity aspect ratio (cavity length/radius). There has been no significant innovation that has broken this linkage without a commensurate sacrifice to the form factor (Footnote 23). The unloaded quality factor, which is limited by the anomalous skin effect [127] for normal metals, changes as

$$Q_0 \propto \nu^{-2/3}. \quad (2.19)$$

It could be improved substantially by switching to partially superconducting cavities, but the presence of the large magnetic field makes benefiting from this approach extremely difficult [109]. The magnetic field itself, despite the promising $R \propto B_0^4$ dependence of Eq. (2.16), is today still within 25–50% of its 1987 value of 6 T [91, 107, 108]. The system noise N_{sys} , however, has improved a few hundredfold between 1987 [97] and 2018 [107].

The improvement to system noise owes to two separate innovations: the use of colder thermal environments culminating in dilution refrigerators to house the haloscopes, and the continued improvement of amplifier technology [104], culminating in Josephson parametric amplifiers capable of adding near the 1/2-photon required by the uncertainty principle [128]. In 2017, HAYSTAC arrived at both of these termini [90, 106] suggesting that a new approach would be needed if noise performance is to continue to improve. This thesis sets out in part to deliver that approach, which will involve circumventing the “quantum limit” on system noise.³¹

Continuing to improve haloscope scan rates is essential in light of the daunting amount of parameter space that remains unscanned. Scan rate optimistically scales with frequency as

$$R \propto \nu^{-14/3} \quad (2.20)$$

for single-mode haloscopes.³² This scaling comes from the cancellation of several frequency-

³¹ There is an analogy to make to Moore’s law, the observation that the number of transistors on computer chips has tended to grow exponentially over time. As transistors approach atomic size scales, the trend cannot possibly continue without a transition to quantum computation. Similarly, it is the neighboring subfield of quantum measurement that may carry axion haloscopes beyond their version of tiny transistors: the uncertainty principle (in reality, most computational tasks do not benefit from quantum computing, so the outlook for axion haloscopes — which only exist to do one task — is in a sense brighter).

³² Haloscopes are furthermore known to achieve 75% of the limit on integrated axion sensitivity set by the Bode-Fano

dependencies of the signal and noise in axion haloscopes, leaving just the dependence of the cavity volume and quality factor on frequency, Eqs.(2.18) and (2.19) [42]. It optimistically assumes $\zeta = 2$ in Eq. (2.18), and neglects the look-elsewhere effect correction to Eq. (2.16) discussed in Sec. 6.5. Equation (2.20) implies that today’s leading haloscopes [91, 107] would take 20 millennia to scan *only* the 1–10 GHz frequency decade at DFSZ sensitivity, generously assuming noise at the quantum limit and 100% live time. To emphasize just how optimistic even this 20 millennia figure is, consider a more pessimistic assessment: using only noise levels achieved before this work, a more realistic live time of 10%, and $\zeta = 3$ in Eq. (2.18), a less optimistic dark matter density of $\rho_a = 0.3 \text{ GeV/cm}^3$ (Footnote 15), and aiming for roughly $g_\gamma^{\text{DFSZ}}/10$ to probe the entire QCD band (yellow region in Fig. 2.4) yields ~ 200 trillion years for the 1–10 GHz decade!

This thesis focuses on two distinct but compatible innovations to speed up detection efforts. One is the first use of the quantum measurement techniques to enhance a search for new particle phenomena. In particular, I use squeezed states of microwave light to avoid the standard quantum limit on noise in an axion haloscope, and speed up the scan rate of the HAYSTAC experiment nearly twofold. The second innovation is the development of a novel Bayesian analysis framework tailored to the problem of dark matter axion detection. I apply this technique to both axion haloscope data and to data in a search for ALPs as much as 17 orders of magnitude lower in mass. Together, these two innovations deliver a roughly threefold scan rate enhancement to the HAYSTAC experiment.³³

criterion [129, 130], making them near-optimal vehicles for single-mode, linear measurements of the axion field [95]. For an abridged discussion of the results in Ref. [95], see Ref. [131].

³³ On its own, a threefold scan rate enhancement is insufficient to solve a 20,000+ year problem. The real utility of the advances discussed in this work is that they naturally invite further innovation in both quantum measurement and scan strategy. The threefold scan rate enhancement achieved herein can additionally be combined with future improvements to live times, magnets [132], and cavities [92, 109, 133]. While 20,000 is a large number of years, $\log_3(20000)$ is not so large a number of comparable scan rate enhancements needed to make the problem quite tractable.

Chapter 3

Josephson parametric amplifiers

Two superconductors sandwiching a thin layer of insulator form a Josephson junction [134, 135], the only dissipationless, nonlinear circuit element available at microwave frequencies [136]. Those three properties have enabled physicists to build nearly lossless parametric amplifiers, such as the Josephson parametric amplifier (JPA), capable of manipulating quantum states of microwave light [137–153]. JPAs and related devices have received particular attention in recent years for their usefulness across a range of quantum optics and quantum information processing applications using superconducting circuits [154–170]. In this thesis we will ultimately use them to improve axion haloscopes beyond the quantum limit [128]. This chapter will explore the physics of JPAs, how they achieve amplification and squeezing, and their relevant design considerations. We will also shed light on one of the core approximations used to model JPAs, a discussion with implications for improving power handling in future designs. We will end by examining the specific JPAs used throughout this thesis.

3.1 The Josephson junction

The Josephson junction, shown schematically in Fig. 3.1 is the element at the heart of many superconducting circuits. It consists of two superconductors separated by an insulating layer (Fig. 3.1a) which Cooper pairs of electrons may tunnel through, forming a current. It is characterized by a critical current I_c , equal to the maximum superconducting current that can flow through it. The critical current that a junction will have when superconducting can be determined

from its room temperature resistance R_s via the Ambegaokar-Baratoff relation [171],

$$I_c = \frac{\pi\Delta(T)}{2eR_n} \tanh\left(\frac{\Delta(T)}{2k_B T}\right), \quad (3.1)$$

where $\Delta(T)$ is the temperature (T)-dependent superconducting energy gap of the material. The dynamics of the junction are determined by the phase drop

$$\delta = \theta_1 - \theta_2 \quad (3.2)$$

of the macroscopic superconducting wavefunction across it. For a more complete introduction to Josephson physics, see Refs. [136, 172].

We will take as a starting point the two equations³⁴ relating this phase drop to the voltage V across and the current I through the junction:

$$I(t) = I_c \sin \delta(t) \quad (3.3)$$

$$V(t) = \frac{\Phi_0}{2\pi} \dot{\delta}(t), \quad (3.4)$$

where $\Phi_0 = h/2e \approx 2.07 \times 10^{-15}$ Wb is the magnetic flux quantum, h is Planck's constant, e is the positron charge, and the dot denotes a time (t) derivative. Instead of working primarily with the voltage and current, we will often find it more useful to refer to their time integrals, the flux and charge variables

$$\Phi(t) = \int_{-\infty}^t V(t') dt' \quad (3.5)$$

$$Q(t) = \int_{-\infty}^t I(t') dt'. \quad (3.6)$$

Equivalently,

$$V(t) = \dot{\Phi}(t) \quad (3.7)$$

$$I(t) = \dot{Q}(t). \quad (3.8)$$

We see from Eqs. (3.4) and (3.7) that the $\Phi = \phi_0 \delta$, where the reduced flux quantum³⁵ $\phi_0 = \Phi_0/2\pi = \hbar/2e$.

³⁴ These are known as the Josephson relations.

³⁵ In superconductivity, the flux through a closed superconducting loop is quantized in units of Φ_0 .




Figure 3.1: (a) A superconductor-insulator-superconductor (SIS) junction forms a Josephson junction. We will treat the Josephson junction as a two-node electrical object. (b) Circuit representation of and conventions used for the Josephson junction. The junction is symbolically represented as an ‘X.’ It is specified by its critical current I_c , which relates the current I flowing through it to the phase drop δ across the junction. The voltage V across the junction is proportional to the time derivative of the phase drop.

3.1.1 Josephson inductance

The Josephson junction behaves as a nonlinear inductor.³⁶ To understand the junction as an inductor, we begin by combining Eqs. (3.3) and (3.4) to relate the voltage to the time derivative of the current;

$$V = \frac{L_{J0}}{\cos \delta} \dot{I}, \quad (3.9)$$

where

$$L_{J0} = \frac{\phi_0}{I_c} \quad (3.10)$$

is a constant describing the junction with dimensions of inductance.

The Josephson inductance can be defined in two different ways: either as the absolute value

³⁶ The Josephson junction can also be modeled as having a self-capacitance, giving rise to a resonance frequency known as the plasma frequency of the junction, and nonzero loss [136]. For our applications, these two effects can be neglected.

of the flux-to-current ratio³⁷

$$L_J = \frac{\Phi}{I} = L_{J0} \frac{|\arcsin(I/I_c)|}{I/I_c}, \quad (3.11)$$

or through Eq. (3.9) as the voltage-to-time derivative of current ratio,

$$L_{J2} = \frac{V}{\dot{I}} = \frac{L_{J0}}{\sqrt{1 - (I/I_c)^2}}. \quad (3.12)$$

Both definitions are valid, and which we use will depend on the context [173].

3.1.2 Josephson energy

An energy can also be associated with the Josephson junction. This is the energy that it takes to build up a certain amount of current flowing through the junction, or equivalently, a certain phase drop across the junction. We know from Eq. (3.3) that applying a current to the junction sets a phase drop across it and from Eq. (3.4) that phase drop's temporal behavior sets a voltage across the junction. Taking energy to be the time integral of power, we have

$$E_J(t) = \int I(t')V(t')dt' = E_{J0}(1 - \cos \delta), \quad (3.13)$$

where

$$E_{J0} = \phi_0 I_c \quad (3.14)$$

has dimensions of energy. We have taken advantage of the fact that energy is only defined up to an additive constant to ensure a minimum value of 0 when no current is flowing.

3.2 The DC SQUID

We have now seen how the Josephson junction acts as a dissipationless, nonlinear inductance. In this section, we will see that a pair of Josephson junctions in parallel maintains all of these essential features, while making the inductance tunable with an externally applied magnetic flux.

³⁷ I will leave out the explicit time dependence in contexts where it is less important to emphasize.

The DC superconducting quantum interference device (SQUID) consists of two Josephson junctions I_{c1} and I_{c2} in parallel, shown schematically in Fig. 3.2.³⁸ In order to derive key properties of the SQUID, it will be useful to work with

$$I = I_1 + I_2 \quad (3.15)$$

$$I_{\text{circ}} = \frac{I_1 - I_2}{2}, \quad (3.16)$$

respectively the total and circulating combinations of the currents I_1 and I_2 through the junctions. The voltage across the SQUID is V , and a magnetic field \mathbf{B}_{tot} may penetrate the plane of the loop as a magnetic flux

$$\Phi_{\text{tot}} = \int_a \mathbf{B}_{\text{tot}} \cdot d\mathbf{a}, \quad (3.17)$$

where a is the area of the loop and $d\mathbf{a}$ is the orthogonal-pointing surface element.

To understand how this object functions as an effective tunable junction, it is necessary to delve into its internal structure. I will treat just the case of two identical junctions

$$I_{c1} = I_{c2} = I_c \quad (3.18)$$

in a symmetrical loop with geometrical inductance L_g , which will be used throughout this work.

3.2.1 Critical current and inductance

The junction phase drops δ_1 and δ_2 are related by the flux penetrating the SQUID loop. To prove this, consider the integration path illustrated in Fig. 3.3 and the four labelled points — one above and below each junction. The magnetic vector potential \mathbf{A} is obtained from the gradient of the phase of the superconducting wave function via the relation [172]

$$\mathbf{A} = -\phi_0 \nabla \theta. \quad (3.19)$$

Therefore,

$$\int_A^B \nabla \theta \cdot d\mathbf{l} + \int_C^D \nabla \theta \cdot d\mathbf{l} = -\frac{1}{\phi_0} \left[\int_A^B \mathbf{A} \cdot d\mathbf{l} + \int_C^D \mathbf{A} \cdot d\mathbf{l} \right]. \quad (3.20)$$

³⁸ The radio frequency (RF) SQUID is the other popular SQUID configuration. It consists of a superconducting loop with a junction in just one arm [172]. This thesis only relies upon the DC SQUID, so henceforth I will drop the DC prefix.




Figure 3.2: Circuit representation of and conventions used for the direct current (DC) superconducting quantum interference device (SQUID). The SQUID consists of two Josephson junctions with critical currents I_{c1} and I_{c2} on opposing arms of a superconducting loop. The direction of positive current flow is given by the arrows. The total current (I , green) and circulating current (I_{circ} , blue) are defined from the currents I_1 (orange) and I_2 (purple) through the junctions, Eq. (3.3). The voltage across the SQUID is V . The magnetic flux Φ_{tot} penetrating the loop arises from any externally applied flux and from I_{circ} in conjunction with the geometrical inductance L_g of the loop.




Figure 3.3: A useful integration path (orange dashed line) for the magnetic vector potential around a SQUID loop. Labels above and below each junction show the integration limits used in Sec. 3.2.1.

Recognizing the sums of integrals on both sides of the equation as the full loop integrals minus the other segments, we have

$$\oint \nabla\theta \cdot d\mathbf{l} - \int_B^C \nabla\theta \cdot d\mathbf{l} - \int_D^A \nabla\theta \cdot d\mathbf{l} = -\frac{1}{\phi_0} \oint \mathbf{A} \cdot d\mathbf{l} + \frac{1}{\phi_0} \left[\int_B^C \mathbf{A} \cdot d\mathbf{l} + \int_D^A \mathbf{A} \cdot d\mathbf{l} \right]. \quad (3.21)$$

The phase at any point on the circuit must be single-valued, so we label the first integral $-2\pi n$, where n is an integer. The next two integrals, by the fundamental theorem of calculus, reduce to the differences of the integration bounds. The first term on the right, using $\mathbf{B} = \nabla \times \mathbf{A}$ with \mathbf{B} being the magnetic field local to the loop, is $-\Phi_{\text{tot}}/\phi_0$. Equation (3.21) now reduces to

$$-(\theta_C - \theta_B) - (\theta_A - \theta_D) - \frac{1}{\phi_0} \left[\int_B^C \mathbf{A} \cdot d\mathbf{l} + \int_D^A \mathbf{A} \cdot d\mathbf{l} \right] = 2\pi n - \frac{\Phi_{\text{tot}}}{\phi_0}. \quad (3.22)$$

In gauge-invariant form, the phase drops are expressed as

$$\delta_1 = \theta_A - \theta_D + \frac{1}{\phi_0} \int_D^A \mathbf{A} \cdot d\mathbf{l} \quad (3.23)$$

$$\delta_2 = \theta_B - \theta_C + \frac{1}{\phi_0} \int_C^B \mathbf{A} \cdot d\mathbf{l}, \quad (3.24)$$

reducing Eq. (3.22) to

$$\delta_2 = \delta_1 - \frac{\Phi_{\text{tot}}}{\phi_0} + 2\pi n. \quad (3.25)$$

This relation will allow us to eliminate δ_2 and work with the enclosed flux Φ_{tot} , a quantity we can influence more readily in the lab.

Using Eq. (3.4) for either junction, we can calculate the inductance of the SQUID. Assuming Φ_{tot} is time-independent, we see that

$$V = \dot{I} \frac{\phi_0 / \left(2I_c \cos \left[\frac{\pi \Phi_{\text{tot}}}{\Phi_0} \right] \right)}{\sqrt{1 - \left(2I_c \cos \left[\frac{\pi \Phi_{\text{tot}}}{\Phi_0} \right] \right)^2}} \quad (3.26)$$

Comparing to Eq. (3.12), we see that the SQUID behaves exactly as a Josephson junction with flux-dependent critical current

$$I_S(\Phi_{\text{tot}}) = 2I_c \left| \cos \left[\frac{\pi \Phi_{\text{tot}}}{\Phi_0} \right] \right|, \quad (3.27)$$

where the absolute value ensures positivity. This result makes intuitive sense: the factor of two owes to having two paths in parallel for the current to take. The effect of an applied flux is to create a circulating current, reducing the amount of total current that can flow through the SQUID before *one* of the junctions reaches its critical current. The SQUID inductance and energy³⁹ are analogous to Eqs. (3.11) and (3.13), respectively:

$$L_{\text{SQ}}(\Phi_{\text{tot}}) = L_S(\Phi_{\text{tot}}) \frac{\arcsin[I/I_S(\Phi_{\text{tot}})]}{I/I_S(\Phi_{\text{tot}})} \quad (3.28)$$

$$E_{\text{SQ}}(t) = E_S(\Phi_{\text{tot}}) \left(1 - \cos \frac{\Phi}{\phi_0}\right), \quad (3.29)$$

where Φ is the flux drop⁴⁰ across the SQUID, Eq. (3.5), and

$$L_S(\Phi_{\text{tot}}) = \frac{\phi_0}{I_S(\Phi_{\text{tot}})} \quad (3.30)$$

$$E_S(\Phi_{\text{tot}}) = \phi_0 I_S(\Phi_{\text{tot}}). \quad (3.31)$$

We also define

$$I_{S0} = 2I_c \quad (3.32)$$

$$L_{S0} = \frac{\phi_0}{I_{S0}} \quad (3.33)$$

$$E_{S0} = \phi_0 I_{S0} \quad (3.34)$$

as the maximum value for the SQUID critical current, and linear terms for inductance and energy when no flux threads the loop, respectively. Before we proceed to treat the SQUID as a Josephson junction with a flux-tunable critical current, we must consider the effect of a nonnegligible geometrical inductance.

³⁹ This is the energy associated with the total current flowing through the SQUID, and does not include any energy stored in the junctions by a circulating current.

⁴⁰ Though they have the same dimensions, the flux drop across (Φ) and penetrating (Φ_{tot}) the SQUID loop are distinct physical quantities. The former is defined as the time integral of the voltage across the SQUID, whereas the latter is a flux in the more familiar sense of a field penetrating an area.

3.2.2 Geometrical inductance and flux-trapping

The effect of nonzero geometrical inductance L_g in the SQUID loop, Fig. 3.2, is to modify the amount of flux threading the loop from the externally applied value Φ_{ext} to a new value⁴¹

$$\Phi_{\text{tot}} = \Phi_{\text{ext}} + L_g I_{\text{circ}}. \quad (3.35)$$

The modification to the external flux will be expressed as a transcendental equation, and give rise to a design rule for SQUIDS whose Φ_{tot} we wish to keep as a single-valued function of Φ_{ext} . Combining Eqs. (3.16) and (3.25), we obtain

$$I_{\text{circ}} = \frac{I_c}{2} \left[\sin \left(\delta_1 - \frac{\Phi_{\text{tot}}}{\phi_0} \right) - \sin \delta_1 \right] = -I_c \sin \left(\pi \frac{\Phi_{\text{tot}}}{\Phi_0} \right) \cos \left(\delta_1 - \pi \frac{\Phi_{\text{tot}}}{\Phi_0} \right). \quad (3.36)$$

Eliminating δ_1 in favor of the total current I , a quantity we can impose experimentally, yields

$$I_{\text{circ}} = -I_c \sin \left(\pi \frac{\Phi_{\text{tot}}}{\Phi_0} \right) \sqrt{1 - \left[\frac{I}{I_S(\Phi_{\text{tot}})} \right]^2}. \quad (3.37)$$

Substituting Eq. (3.37) into Eq. (3.35) gives us an expression relating the applied and total fluxes:

$$\Phi_{\text{tot}} = \Phi_{\text{ext}} - L_g I_c \sin \left(\pi \frac{\Phi_{\text{tot}}}{\Phi_0} \right) \sqrt{1 - \left[\frac{I}{I_S(\Phi_{\text{tot}})} \right]^2}. \quad (3.38)$$

This is our transcendental equation.

When

$$\frac{\partial \Phi_{\text{ext}}}{\partial \Phi_{\text{tot}}} < 0, \quad (3.39)$$

Φ_{tot} takes on multiple values for a given Φ_{ext} , as illustrated in Fig. 3.4. Here, the SQUID becomes hysteretic, and enters the “flux-trapping” regime, which is undesirable for controlling it. To quantify at what value of L_g flux-trapping sets in, we will consider the simple case where $I = 0$ ⁴² so that

$$\Phi_{\text{ext}} = \Phi_{\text{tot}} + L_g I_c \sin \left(\pi \frac{\Phi_{\text{tot}}}{\Phi_0} \right), \quad (3.40)$$

⁴¹ For the symmetric SQUID loop, Eq. (3.18), total current I will not directly produce any additional magnetic flux, as the two branches of the SQUID will contribute flux equally and oppositely.

⁴² Without this assumption the math gets considerably more opaque. The condition we derive with $I = 0$ will suffice for us as a design rule.




Figure 3.4: Total versus externally applied flux for a SQUID with different geometrical inductances, Eq. (3.42). When $\beta_L > 2$ (purple and brown curves) some values of Φ_{ext} become associated with multiple values of Φ_{tot} , and the SQUID becomes hysteretic.

which differentiates to

$$\frac{\partial \Phi_{\text{ext}}}{\partial \Phi_{\text{tot}}} = 1 + \frac{\beta_L}{2} \cos\left(\frac{\pi \Phi_{\text{tot}}}{\Phi_0}\right), \quad (3.41)$$

where

$$\beta_L = \frac{L_g}{L_{J0}} \quad (3.42)$$

dimensionlessly parametrizes the geometrical inductance. To avoid flux-trapping, Eq. (3.39), we require

$$\beta_L \leq 2. \quad (3.43)$$

So long as we adhere to Eq. (3.43), a SQUID functions as a flux-tunable Josephson junction, which in turn acts like a nonlinear inductor. In order to build circuits out of SQUIDs that can process quantum mechanical signals, we will next need to develop the language of circuit quantization.

3.3 Circuit quantization

The simple harmonic oscillator (SHO) is a canonical introductory example for students new to quantum mechanics [174]. Though it is often introduced in a mechanical context — as a mass on a spring or hanging from a pendulum — the physics applies just as well to electrical oscillators. We will first study the instructive example of a simple resonant circuit, and then apply the technique to a circuit which incorporates the SQUIDS of Sec. 3.2 in order to perform parametric amplification.

3.3.1 LC circuit

The circuit of Fig. 3.5 consists of an inductor L and capacitor C , which alternately store magnetic and electric energy, respectively, as

$$E_L = \frac{1}{2} \frac{\Phi^2}{L} \quad (3.44)$$

$$E_C = \frac{1}{2} C \dot{\Phi}^2. \quad (3.45)$$

We have written everything in terms of a single generalized coordinate: the flux across the elements Φ , Eq. (3.5), and its time derivative. The inductively and capacitively stored energies will be analogous to the potential and kinetic energies of a mechanical oscillator.⁴³




Figure 3.5: Simple LC circuit. The inductor L stores magnetic energy and the capacitor C stores electric energy. Together, they form a resonator as the energy sloshes from one circuit element to the other.

⁴³ In the analogy to a mass m on a spring k , the capacitance C corresponds to m , and the inductance L corresponds to $1/k$.

3.3.2 Lagrangian and Hamiltonian

By writing down the Lagrangian of the system, we will be able to obtain the Hamiltonian, which compactly stores the information on the time evolution of the system. From Eqs. (3.44) and (3.45), we can write down the Lagrangian of the circuit in Fig. 3.5:

$$\mathcal{L} = \frac{1}{2}C\dot{\Phi}^2 - \frac{1}{2}\frac{\Phi^2}{L}. \quad (3.46)$$

The momentum conjugate to Φ is obtained as

$$\frac{\partial \mathcal{L}}{\partial \dot{\Phi}} = C\dot{\Phi} = Q. \quad (3.47)$$

Physically, it is the charge on one capacitor plate, Eq. (3.6). To obtain the Hamiltonian from the Lagrangian, we perform a Legendre transformation:

$$\mathcal{H}(\Phi, Q) = \dot{\Phi} \frac{\partial \mathcal{L}}{\partial \dot{\Phi}} - \mathcal{L} = \frac{1}{2} \frac{Q^2}{C} + \frac{1}{2} \frac{\Phi^2}{L}, \quad (3.48)$$

which is the sum of the circuit element energies.

3.3.3 Classical equations of motion

Before quantizing the system, we can check our understanding by writing down the Hamilton's equations of motion from Eq. (3.48):

$$\frac{d\Phi}{dt} = \frac{\partial \mathcal{H}}{\partial Q} = \frac{Q}{C} \quad (3.49)$$

$$\frac{dQ}{dt} = -\frac{\partial \mathcal{H}}{\partial \Phi} = -\frac{\Phi}{L} \quad (3.50)$$

The first equation, using Eq. (3.5), is simply the constitutive equation for the capacitor, commonly written

$$C = \frac{Q}{V}. \quad (3.51)$$

Using the definition for the conjugate momentum, Eq. (3.47), the second equation yields the equation of motion for a SHO,

$$C\ddot{\Phi} = -\frac{\Phi}{L}, \quad (3.52)$$

which resonates at

$$\omega_{LC} = \frac{1}{\sqrt{LC}}. \quad (3.53)$$

3.3.4 Quantization and ladder operators

By analogy with the position and momentum of a mechanical oscillator, we promote Φ and Q to Hermitian operators $\hat{\Phi}$ and \hat{Q} , giving a Hamiltonian operator

$$\hat{H} = \frac{1}{2} \frac{\hat{Q}^2}{C} + \frac{1}{2} \frac{\hat{\Phi}^2}{L}. \quad (3.54)$$

These new operators satisfy the canonical commutation relationship

$$[\hat{\Phi}, \hat{Q}] = i\hbar, \quad (3.55)$$

and can be expressed in terms of the ladder operators⁴⁴

$$\hat{A} = \frac{1}{\sqrt{2\hbar Z_{LC}}} (\hat{\Phi} + iZ_{LC}\hat{Q}) \quad (3.56)$$

$$\hat{A}^\dagger = \frac{1}{\sqrt{2\hbar Z_{LC}}} (\hat{\Phi} - iZ_{LC}\hat{Q}), \quad (3.57)$$

where

$$Z_{LC} = \sqrt{\frac{L}{C}} \quad (3.58)$$

is the impedance of the resonator and

$$[\hat{A}, \hat{A}^\dagger] = 1 \quad (3.59)$$

follows from Eq. (3.55). The inverse relationships to Eqs. (3.56) and (3.57) are

$$\hat{\Phi} = \sqrt{\frac{\hbar Z_{LC}}{2}} (\hat{A}^\dagger + \hat{A}) \quad (3.60)$$

$$\hat{Q} = i\sqrt{\frac{\hbar}{2Z_{LC}}} (\hat{A}^\dagger - \hat{A}). \quad (3.61)$$

In terms of the ladder operators, the Hamiltonian operator, Eq. (3.54), takes the familiar form

$$\hat{H} = \hbar\omega_{LC} \left(\hat{A}^\dagger \hat{A} + \frac{1}{2} \right) = \hbar\omega_{LC} \left(\hat{n} + \frac{1}{2} \right), \quad (3.62)$$

⁴⁴ The ladder operators are commonly referred to as the creation, or raising, operator A^\dagger , and the annihilation, or lowering, operator A for how they act upon energy eigenstates of the Hamiltonian [174].

where $\hat{n} = \hat{A}^\dagger \hat{A}$ is the number operator, which acts on an energy eigenstate $|n\rangle$ as

$$\hat{n} |n\rangle = n |n\rangle \quad (3.63)$$

to return the number n of photons with energy $\hbar\omega_{LC}$ occupying the resonator. The $+1/2$ -photon term in Eq. (3.62) is the ground state energy of the oscillator.⁴⁵

3.3.5 Heisenberg equation of motion

The Hamiltonian, Eq. (3.54) or Eq. (3.62), is nothing more than a suitcase for compactly carrying around the time-evolution behavior of a nondissipative system such as our LC circuit. To unpack the time dependence, we use the Heisenberg equation of motion for the time evolution of an operator \hat{O} [175]:⁴⁶

$$\frac{d\hat{O}}{dt} = \frac{i}{\hbar} [\hat{H}, \hat{O}]. \quad (3.64)$$

For the example of our simple LC circuit, this yields

$$\frac{d\hat{A}}{dt} = -i\omega_{LC}\hat{A}. \quad (3.65)$$

If \hat{A} were a complex field amplitude, Eq. (3.65) says that it would simply evolve in phase without changing amplitude, as is expected for an undriven, undamped resonator.

This concludes our introductory discussion of circuit quantization. We will take the language of this section and apply it to more complex devices made out of the components of this and previous sections to model their behavior.

3.4 Lumped-element Josephson parametric amplifier

Parametric amplification is a process that will allow us to both amplify (as the name suggests) and also deamplify quantum observables [137, 176]. It is achieved by modulating a resonator such as that modeled in Fig. 3.6 at approximately twice its resonance frequency. Equation (3.28), which

⁴⁵ The $+1/2$ will correspond to the quantum noise in systems we are interested in measuring.

⁴⁶ This equation works in the Heisenberg picture of quantum mechanics, in which operators evolve. It assumes that \hat{O} contains no explicit time dependence in the Schrödinger picture.

Taylor expands in the $I/I_S \ll 1$ limit to

$$L_{\text{SQ}}(\Phi_{\text{tot}}) = L_S(\Phi_{\text{tot}}) \left[1 + \frac{1}{6} \left(\frac{I}{I_S(\Phi_{\text{tot}})} \right)^2 + \mathcal{O} \left(\left[\frac{I}{I_S(\Phi_{\text{tot}})} \right]^4 \right) \right] \quad (3.66)$$

suggests two possible ways of doing this with a microwave resonator containing SQUIDs. First, the leading I^2 -dependence of the inductance provides a natural frequency-doubling: a current oscillating at frequency ω will vary the inductance at 2ω . Second, the flux-dependent SQUID inductance $L_S(\Phi_{\text{tot}})$ can be directly modulated at 2ω . These two options correspond to applying a single, monochromatic current-pump tone (henceforth, current pumping)⁴⁷ [180, 181] or a single, monochromatic magnetic flux-pump tone (henceforth, flux pumping) [141, 182–186] to the SQUID(s) in the resonator.

The resonators we will work with exist in the lumped-element regime [187], wherein their spatial extent is much smaller than a wavelength of light at any frequency of interest.⁴⁸ In particular, the $4 \text{ GHz} < \nu < 15 \text{ GHz}$ frequency window that this thesis operates within corresponds to wavelengths

$$\lambda = \frac{v_{\text{TL}}}{\nu} \quad (3.67)$$

of 1.3–5.0 cm, where $v_{\text{TL}} \approx 2c/3$ is the approximate speed of light in our transmission line geometries. Our devices, for comparison, will be sub-mm in length.

Figure 3.6a shows the resonator we will be working with: a lumped-element tunable Kerr circuit (TKC), which lies at the heart of our JPAs [188]. In this thesis, “JPA” refers to a two-port device that consists of a TKC plus (at least) one microwave circulator to separate incoming and outgoing signals, as shown in Fig. 3.7. It is common in the literature, however, to refer to the JPA as the one-port reflection device. At the center of the TKC is a resonator formed out of a capacitor C and an inductive series array of N identical, balanced SQUIDs each with critical

⁴⁷ There is a second way to pump using current: by applying two pump tones roughly equally detuned about the central frequency of the resonator [177, 178]. Both of these current-pumping methods correspond to a four-wave mixing process, whereas flux pumping uses three-wave mixing. In this thesis, current pumping will just refer to the use of a single tone. For a comparison of some aspects of all three pumping schemes, see Ref. [179].

⁴⁸ It is common for JPAs to use both lumped-element [143, 145, 147, 188] and distributed [139, 141, 148] resonator designs.




Figure 3.6: (a) Circuit schematic for a tunable Kerr circuit (TKC). The TKC (within the dashed orange box) consists of a series array of N SQUIDs with critical currents I_S in series with a geometrical inductance L_G , all shunted by a capacitance C and coupled to the external environment via a smaller coupling capacitor C_c . The TKC is here shown biased by a Thévenin equivalent voltage source $V e^{-i\omega t}$ with impedance Z_0 . (b) Circuit symbol used to represent the TKC in larger schematics. The TKC will be used as an amplifier in reflection. (c) Input-output theory model of the TKC as a resonant mode \hat{A} (blue) with a resonance ω_{TKC} and Kerr nonlinearity K . The mode couples at power decay rate γ_e to input and output propagating modes \hat{A}_{in} and \hat{A}_{out} , which are decomposed into classical pump (α) and quantum signal (\hat{a}) components.

current $I_S(\Phi_{\text{tot}})$.⁴⁹ The SQUIDs are modeled in series with any geometrical inductance L_G .⁵⁰ Finally, the TKC includes a smaller coupling capacitor C_c in order to control the coupling rate to the external environment. Here, that environment is represented as a Thévenin equivalent voltage source applying a stimulus at some frequency ω through an impedance Z_0 . For us, $Z_0 = 50 \Omega$, standard for microwave equipment. Starting in Sec. 3.4.1, we will first derive some basic properties of the TKC by analyzing the circuit diagram. In Sec. 3.4.2 we will obtain the Hamiltonian, which we approximate as that of a Duffing (or Kerr) oscillator in Sec. 3.4.3. Finally, in Secs. 3.4.4–3.4.7 we will use input-output theory (IOT) [189, 190] to understand its capabilities for amplification, deamplification, and squeezing.

⁴⁹ We are assumed not to be in the flux-trapping regime of Sec. 3.2.2, such that a magnetic field Φ_{ext} applied uniformly across the SQUIDs gives rise to the same field Φ_{tot} in each loop via Eq. (3.35).

⁵⁰ The SQUID loop geometrical inductance L_g of Eq. (3.35) contributes to the TKC series inductance L_G .




Figure 3.7: Relationship between TKC (Fig. 3.6b) and Josephson parametric amplifier (JPA; triangle). The JPA is represented with the usual triangular symbol for amplifiers with input (output) port on the left (right). Inside (in gray), it consists of a microwave circulator, which routes propagating waves in the sense indicated by the arrow on it, and the TKC, Fig. 3.6b, which processes signals that reflect off of it. Not shown is the line used to provide the pump tone, as the implementation differs between current- and flux-pumped designs.

3.4.1 Basic TKC circuit analysis

Basic properties of the circuit — its resonance frequency and external quality factor — can be obtained from considering the impedances of and energies stored or dissipated within the circuit elements. Here, we will treat the series array of SQUIDs, whose *linear* inductance components add like simple linear inductors, as a single inductance

$$L_a = NL_S, \quad (3.68)$$

where L_S comes from Eq. (3.28) and I have dropped the dependence on Φ_{tot} for notational convenience. We further combine the inductive array with the geometrical inductance as

$$L_{aG} = L_a + L_G. \quad (3.69)$$

Since the coupling capacitor is in series with a reactive combination of circuit elements (Fig. 3.6a), resonance will occur when the series impedance⁵¹

$$Z_{\text{series}} = Z_{C_c} + Z_{L_{aG}C} \quad (3.70)$$

⁵¹ Throughout this section, we assume a weakly coupled resonator $C_c \ll C$, with the absolute values of the impedances of L and C near resonance roughly equal to Z_0 . This is necessary to use the condition for a series resonance.

vanishes, where

$$Z_{L_{aG}C} = \frac{L_{aG}/C}{i\omega L_{aG} + \frac{1}{i\omega C}} \quad (3.71)$$

$$Z_{C_c} = \frac{1}{i\omega C_c}. \quad (3.72)$$

This occurs at

$$\omega_{\text{TKC}} = \frac{1}{\sqrt{L_{aG}C_{\Sigma}}}, \quad (3.73)$$

where

$$C_{\Sigma} = C + C_c. \quad (3.74)$$

The external quality factor⁵² Q_e quantifies what amount of the energy stored in the resonator decays out to the external environment Z_0 with each period of oscillation as

$$Q_e = \omega_{\text{TKC}} \frac{E_{\text{stored}}}{P_{\text{out}}}, \quad (3.75)$$

where E_{stored} and P_{out} are the energy stored in and power decaying out of the resonator on resonance. For the TKC, using Eqs. (3.44) and (3.45) for the energy stored in inductors and capacitors,

$$Q_e \approx \frac{\sqrt{L/C}}{Z_0} \left(\frac{C}{C_c} \right)^2 \approx \left(\frac{C}{C_c} \right)^2. \quad (3.76)$$

To obtain the first equality, we have used the approximation $C \gg C_c$, which is valid for the JPAs discussed in this thesis. In practice, we will design $Z_{LC} \approx Z_0$, as trying to push the inductor or capacitor impedance magnitude far from this value can give rise to parasitic reactances, causing them to no longer act as intended. The second equality therefore provides a reasonable approximation for the external quality factor.

Finally, the external power decay rate of the resonator

$$\gamma_e = \frac{\omega_{\text{TKC}}}{Q_e} \approx \frac{Z_0}{L} \left(\frac{C_c}{C} \right)^2, \quad (3.77)$$

corresponds to the 2π times the width (in Hz) of the Lorentzian response profile of the resonator.

With these useful quantities in hand, we will now obtain the Hamiltonian in order to study the nonlinearity and the dynamics of the system.

⁵² For our purposes, it will be a reasonable approximation to model the TKC as having no internal loss, so the internal quality factor is treated as infinite. A simple way to model internal loss would be with a resistor in parallel with the capacitor and array-plus-inductor.

3.4.2 TKC Lagrangian and Hamiltonian

The quantization procedure of Sec. 3.3 requires us to write down the energies of the elements in Fig. 3.6a in terms of a generalized coordinate and its derivative. For simplicity, we will set $L_G = 0$ and return to its significance in Sec. 3.5 in relation to SQUID arrays. We will also neglect the voltage source and its impedance, and add in the external coupling rate of Sec. 3.4.1 by hand in Sec. 3.4.4. Our coordinate will be the flux drop Φ_a across the array, which is the sum of the N identical flux drops Φ_{SQ} across the individual SQUIDs:⁵³

$$\Phi_a = N\Phi_{\text{SQ}}. \quad (3.78)$$

The magnetic energy E_m associated with the array, using Eq. (3.29), is therefore

$$E_m = E_a \left(1 - \cos \frac{\Phi_a}{N\phi_0} \right), \quad (3.79)$$

where

$$E_a = NE_S \quad (3.80)$$

via Eq. (3.31).

The electric energy E_e associated with the capacitors is given for each capacitor by Eq. (3.45). Kirchhoff's voltage law requires that voltage across C_c be equal and opposite to that across C , making $\dot{\Phi}_{C_c}^2 = \dot{\Phi}_C^2 = \dot{\Phi}_a^2$ and

$$E_e = \frac{1}{2}C_\Sigma \dot{\Phi}_a^2. \quad (3.81)$$

We can now obtain the Hamiltonian of the TKC by first obtaining the Lagrangian:

$$\mathcal{L}_{\text{TKC}}(\Phi_a, \dot{\Phi}_a) = \frac{1}{2}C_\Sigma \dot{\Phi}_a^2 - E_a \left(1 - \cos \frac{\Phi_a}{N\phi_0} \right). \quad (3.82)$$

The momentum conjugate to Φ_a is given by

$$\frac{\partial \mathcal{L}_{\text{TKC}}}{\partial \dot{\Phi}_a} = C_\Sigma \dot{\Phi}_a = CV_C - C_c V_{C_c} = Q_\Delta, \quad (3.83)$$

⁵³ We are here assuming that the SQUID array has no internal dynamics of interest. If the SQUIDs are nearly identical, are embedded in similar surroundings to one another, are not in the flux-trapping regime, and begin with the same initial conditions, this is a reasonable assumption.

where the second equality uses Eq. (3.7) and the last equality interprets the quantity via Eq. (3.51) as the difference in charge on the “positive” plates of the capacitors as denoted in Fig. 3.6a. Performing the Legendre transformation of Eq. (3.48) yields the classical Hamiltonian,

$$\mathcal{H}_{\text{TKC}}(\Phi_a, Q_\Delta) = \frac{1}{2} \frac{Q_\Delta^2}{C_\Sigma} + E_a \left(1 - \cos \frac{\Phi_a}{N\phi_0} \right). \quad (3.84)$$

3.4.3 Kerr Hamiltonian

Rather than dealing with the full cosine dependence of the Hamiltonian, Eq. (3.84), we can Taylor expand for $\Phi_a/N\phi_0 \ll 1$, equivalent to $\Phi_{\text{SQ}} \ll \phi_0$ for the flux drop across each SQUID or $I \ll I_c$ for the current through each junction. Performing this expansion, we obtain

$$\mathcal{H}_{\text{TKC}}(\Phi_a, Q_\Delta) \approx \frac{Q_\Delta^2}{2C_\Sigma} + \frac{E_a}{2} \left(\left[\frac{\Phi_a}{N\phi_0} \right]^2 - \frac{1}{12} \left[\frac{\Phi_a}{N\phi_0} \right]^4 \right), \quad (3.85)$$

where we have kept only the first term beyond the bare Hamiltonian:

$$\mathcal{H}_0(\Phi_a, Q_\Delta) = \frac{Q_\Delta^2}{2C_\Sigma} + \frac{E_a}{2} \left[\frac{\Phi_a}{N\phi_0} \right]^2. \quad (3.86)$$

Quadratic terms (Φ_a^2, Q_Δ^2) in the Hamiltonian become linear in the equations of motion, and comparing to Eq. (3.48) shows that \mathcal{H}_0 is indeed the Hamiltonian for a linear LC resonator, with $C \rightarrow C_\Sigma$ and

$$L \rightarrow L_a = N \frac{\phi_0}{I_s}, \quad (3.87)$$

exactly as we would expect for N linear inductances given by Eq. (3.30) in series.

Equation (3.85) will prove to be a Kerr Hamiltonian [191–194] of the form

$$\mathcal{H}_K = \hbar\omega_0 A^* A + \hbar \frac{K}{2} A^* A^* A A, \quad (3.88)$$

where A and its complex conjugate A^* are the classical analogues⁵⁴ of the ladder operators A and A^\dagger such that the number of photons worth of energy in the resonator is

$$n = |A|^2. \quad (3.89)$$

⁵⁴ By this point, we could have quantized the system as in Sec. 3.3, but the Kerr constant can be calculated classically, with quantum mechanics providing an obvious motivation for factors of \hbar that appear.

The Kerr Hamiltonian encodes the dynamics of the Duffing oscillator [195–197], with quartic terms in the Hamiltonian corresponding to cubic terms in the equation of motion. The Kerr constant K effectively shifts the resonance of a Duffing oscillator by $K/2$ for every quantum of energy in the resonator. If K is positive (negative), the resonance shifts up (down), and in a mechanical system the effect can be thought of as a nonlinear “hardening” (“softening”) of the spring constant. Following Sec. 3.3.4, we can write down the expressions for A and A^* :

$$A = \frac{1}{\sqrt{2\hbar Z_{\text{TKC}}}} (\Phi_a + iZ_{\text{TKC}}Q_\Delta) \quad (3.90)$$

$$A^* = \frac{1}{\sqrt{2\hbar Z_{\text{TKC}}}} (\Phi_a - iZ_{\text{TKC}}Q_\Delta), \quad (3.91)$$

where

$$Z_{\text{TKC}} = \sqrt{\frac{N\phi_0}{C_\Sigma I_S}} \quad (3.92)$$

has dimensions of impedance and the inverse relations are

$$\Phi_a = \sqrt{\frac{\hbar Z_{\text{TKC}}}{2}} (A^* + A) \quad (3.93)$$

$$Q_\Delta = i\sqrt{\frac{\hbar}{2Z_{\text{TKC}}}} (A^* - A). \quad (3.94)$$

Plugging Eqs. (3.93) and (3.94) into Eq. (3.85),

$$\mathcal{H}_{\text{TKC}} \approx \hbar\omega_{\text{TKC}}n - \frac{1}{16} \frac{\hbar^2}{N^2\phi_0^2 C_\Sigma} \left[n^2 + \frac{2}{3}n(A^2 + A^{*2}) + \frac{1}{6}(A^4 + A^{*4}) \right]. \quad (3.95)$$

The rotating wave approximation (RWA) allows us to ignore terms with unequal factors of A and A^* on the grounds that for dynamics about resonance, they will oscillate rapidly, averaging to zero.

We identify our resonator’s Kerr constant from Eqs. (3.88) and (3.95) as

$$K = -\frac{\hbar}{8N^2\phi_0^2 C_\Sigma} = -\frac{\hbar\omega_{\text{TKC}}^2}{8N\phi_0 I_S}. \quad (3.96)$$

The negative sign indicates that the array nonlinearity acts as a soft spring, pulling resonance down. The second expression is more useful from a design standpoint. It illustrates how K may be changed while keeping ω_{TKC} constant. For example, we could increase both N and I_c (and hence I_S) by some factor m , which would decrease K as m^2 without changing ω_{TKC} . Evidently, adding

more SQUIDS decreases the nonlinearity: physically, the flux drop across each individual SQUID gets linearly smaller with N while the array energy, Eq. (3.79), does not depend on the flux drop to linear order [198].

We have seen that the TKC circuit of Fig. 3.6a can behave as a Duffing oscillator, or equivalently, a resonator with a Kerr nonlinearity. In doing so, we have connected the parameters ω_{TKC} , Q_e , γ_e , and K , which determine the behavior of the device, to the design parameters I_c , N , C , C_c , and L_G and the experimental knob Φ_{tot} . Next, we will see one way in which such a device can deliver parametric gain.

3.4.4 Input-output theory derivation of gain

JPAAs amplify small signals at their input by “stealing” some of the power from a much larger⁵⁵ pump tone. Here, we will derive the gain, deamplification, and squeezing of a current-pumped JPA using IOT. The goal of IOT, as the name suggests, is to write down equations for the output fields in terms of the input fields which do not include the internal degrees of freedom of the system. This derivation closely follows that of Ref. [148]. For a more detailed derivation, see Ref. [180].

We take as our starting point the IOT model of Fig. 3.6c. The cavity is modelled with the quantized Kerr Hamiltonian

$$\hat{H}_K = \hbar\omega_0\hat{A}^*\hat{A} + \hbar\frac{K}{2}\hat{A}^\dagger\hat{A}^\dagger\hat{A}\hat{A}, \quad (3.97)$$

whose classical counterpart is Eq. (3.88). Input and output fields⁵⁶

$$\hat{A}_{\text{in}} = (\alpha_{\text{in}} + \hat{a}_{\text{in}})e^{-i\omega_p t} \quad (3.98)$$

$$\hat{A}_{\text{out}} = (\alpha_{\text{out}} + \hat{a}_{\text{out}})e^{-i\omega_p t}, \quad (3.99)$$

where α_{in} and α_{out} are respectively the incoming and outgoing strengths of a classical pump field at frequency ω_p and \hat{a}_{in} and \hat{a}_{out} are the corresponding quantum signals,⁵⁷ are coupled to the cavity

⁵⁵ As a rule of thumb for maintaining a sufficiently “stiff” pump tone for current pumping in practice, that pump tone should be at least 20 dB larger than the *amplified* signal [188].

⁵⁶ All fields are for now functions of time t , where I will leave the functional dependence out for notational convenience where it is not needed.

⁵⁷ The pump can be treated as classical because the effect of its quantum fluctuations within the resonator can be made much smaller than those of the signal by weakly coupling in a large tone to the shared signal/pump line.

at power decay rate⁵⁸ γ_e , Eq. (3.77). Factoring out the trivial time dependence $e^{-i\omega_p t}$ allows us to work in the rotating frame of the pump. The input and output fields obey the boundary condition

$$\sqrt{\gamma_e}\hat{A} = \hat{A}_{\text{in}} - \hat{A}_{\text{out}}, \quad (3.100)$$

and standard bosonic commutation relations subject to a Markov approximation [137, 189]. The cavity mode \hat{A} is decomposed into pump (α) and signal (\hat{a}) components as in Eqs. (3.98) and (3.99):

$$\hat{A} = (\alpha + \hat{a})e^{-i\omega_p t}. \quad (3.101)$$

We follow the standard IOT approach [189, 190]. Using Eq. (3.64), we will write down the Heisenberg equation of motion for \hat{A} in Eq. (3.97). The Heisenberg equation of motion only models the closed system, however. To model the interaction with the incoming and outgoing fields, we must add on the terms modeling dissipation and associated fluctuation by hand, forming the Heisenberg-Langevin equation of motion. The Heisenberg-Langevin equation is then solved in the steady state for large, classical input and output fields α_{in} and α_{out} corresponding to our pump tone. We use the solution to linearize around the small input and output fields \hat{a}_{in} and \hat{a}_{out} , ending with an equation for \hat{a}_{out} in terms of only input fields.

Applying Eq. (3.64) to Eq. (3.97) gives the Heisenberg equation for the internal dynamics of the system:

$$\frac{d\hat{A}}{dt} = -i\omega_{\text{TKC}}\hat{A} - iK\hat{A}^\dagger\hat{A}\hat{A}. \quad (3.102)$$

Inserting the coupling terms by hand gives the Heisenberg-Langevin equation

$$\frac{d\hat{A}}{dt} = -i\omega_{\text{TKC}}\hat{A} - iK\hat{A}^\dagger\hat{A}\hat{A} - \frac{\gamma_e}{2}\hat{A} + \sqrt{\gamma_e}\hat{A}_{\text{in}}, \quad (3.103)$$

which, together with Eq. (3.100), governs the full open system.

We first solve for the response to a large, classical tone

$$\hat{A}_{\text{in}} \rightarrow \alpha_{\text{in}}e^{-i\omega_p t}. \quad (3.104)$$

⁵⁸ The power decay rate is the inverse of the characteristic time for the energy in the resonator to deplete. Since coupling is reciprocal, it also parametrizes how quickly an external drive adds energy to the cavity. It is equal to twice the field decay rate, and corresponds to the full-width at half-max of the Lorentzian response profile of the cavity. The field decay rate therefore corresponds to the half-width at half-max.

We are interested in the steady-state solution, where

$$\frac{d\alpha}{dt} = 0. \quad (3.105)$$

Plugging Eq. (3.104) and its analogues for α and α_{out} into Eq. (3.102) subject to Eq. (3.105) yields

$$\sqrt{\gamma_e}\alpha_{\text{in}} = i(\omega_{\text{TKC}} - \omega_p)\alpha + iK|\alpha|^2\alpha + \frac{\gamma_e}{2}\alpha. \quad (3.106)$$

Multiplying Eq. (3.106) by its complex conjugate gives

$$\kappa n_{\alpha,\text{in}}^2 = \left(\Delta^2 + \frac{\gamma_e}{2}\right)n_\alpha - 2\Delta K n_\alpha^2 + K^2 n_\alpha^3, \quad (3.107)$$

where

$$\Delta = \omega_p - \omega_{\text{TKC}} \quad (3.108)$$

$$n_{\alpha,\text{in}} = |\alpha_{\text{in}}|^2 \quad (3.109)$$

$$n_\alpha = |\alpha|^2 \quad (3.110)$$

are respectively the detuning of of the pump from the bare resonance, the incoming number of pump photons, and the number of photons from the pump in the resonator. Equation (3.106) will be of considerable use to us in Ch. 4, without need to even directly incorporate the quantum mechanical signal operators.

For now, we recast Eq. (3.107) in terms of the dimensionless quantities

$$\delta = \frac{\Delta}{\gamma_e} \quad (3.111)$$

$$\xi = \frac{|\tilde{\alpha}_{\text{in}}|^2 K}{\gamma_e} \quad (3.112)$$

$$\tilde{n} = \frac{n_\alpha}{|\tilde{\alpha}_{\text{in}}|^2} \quad (3.113)$$

as

$$\xi^2 \tilde{n}^3 - 2\delta\xi \tilde{n}^2 + \left(\delta^2 + \frac{1}{4}\right)\tilde{n} - 1 = 0, \quad (3.114)$$

where

$$\tilde{\alpha}_{\text{in}} = \frac{\alpha_{\text{in}}}{\sqrt{\gamma_e}} \quad (3.115)$$

In Eqs. (3.111), (3.112), (3.113), and (3.115), δ is the dimensionless drive detuning from resonance,⁵⁹ ξ is the dimensionless nonlinearity (or dimensionless drive power, depending on what one is interested in), \tilde{n} is the dimensionless resonator occupancy, and $\tilde{\alpha}_{\text{in}}$ is the dimensionless drive amplitude.

The cubic Eq. (3.114) specifies a critical point $\{\delta_c, \xi_c\}$ where

$$\frac{\partial \delta}{\partial \tilde{n}} = \frac{\partial^2 \delta}{\partial \tilde{n}^2} = 0. \quad (3.116)$$

The critical drive detuning and power are given by

$$\delta_c = -\frac{\sqrt{3}}{2} \quad (3.117)$$

$$\xi_c = -\frac{1}{3\sqrt{3}}, \quad (3.118)$$

and the value of \tilde{n} at the critical point is

$$\tilde{n}_c = 3. \quad (3.119)$$

For a portion of operation space below δ_c and above ξ_c there are three real solutions to \tilde{n} . Two of these — the highest and lowest values for \tilde{n} — will be stable solutions, and the third unstable. Figure 3.8 shows the response of the Duffing oscillator to sub- and super-critical drives. While the bifurcation physics (magenta and black curves in Fig. 3.8) can also be used to read out feeble signals by, for example, correlating the state of a qubit with the high/low amplitude stable states of the oscillator [138], we will not be interested in it here.⁶⁰ Instead, we will work just short of the critical point (typically between the green and light blue curves in Fig. 3.8), where the nonlinear response of the JPA is sharp but single-valued.

With the response to a large pump tone in hand, we now solve for the response to the full input field Eq. (3.101). Equation (3.102) gives

$$\frac{d\hat{a}}{dt} = \left(i\Delta - 2iKn_\alpha + \frac{\gamma_e}{2} \right) \hat{a} - iKn_\alpha \hat{a}^\dagger + \sqrt{\gamma_e} \hat{a}_{\text{in}}, \quad (3.120)$$

⁵⁹ This is an entirely separate quantity from the phase drop δ defined in Eq. (3.2) at the beginning of this chapter.

⁶⁰ The flux-pumped JPAs of Sec. 3.7 likewise have useful bifurcation physics. Reference [170] uses a closely-related design to measure a superconducting qubit without the need for a ferrite circulator prior to amplification.




Figure 3.8: Duffing oscillator response showing internal amplitude α relative to its value at the critical point α_c versus pump frequency ω_p (absent any signal tone) relative to TKC resonant frequency ω_{TKC} . The different curves correspond to different drive powers linearly spaced from $\xi = 0.25\xi_c$ (blue) to $\xi = 1.5\xi_c$ (black). Beyond the critical point, the response becomes multi-valued. Figure adapted from Ref. [173].

where we have thrown out terms higher than linear order in \hat{a} . Using the convention

$$b(t) = \sqrt{\frac{\gamma_e}{2\pi}} \int_{-\infty}^{\infty} d\tilde{\Delta} e^{-i\tilde{\Delta}\gamma_e t} b(\tilde{\Delta}) \quad (3.121)$$

to examine the Fourier domain behavior in terms of the dimensionless detuning

$$\tilde{\Delta} = \frac{\omega - \omega_p}{\gamma_e}, \quad (3.122)$$

Eq. (3.120) becomes

$$\begin{bmatrix} \hat{a}_{\text{in}}(\tilde{\Delta}) \\ \hat{a}_{\text{in}}^\dagger(-\tilde{\Delta}) \end{bmatrix} = \sqrt{\gamma_e} \begin{bmatrix} i(-\delta + 2\xi\tilde{n} - \tilde{\Delta}) + \frac{1}{2} & \xi\tilde{n}e^{2i\phi} \\ -i\xi\tilde{n}e^{-2i\phi} & i(\delta - 2\xi\tilde{n} - \tilde{\Delta}) + \frac{1}{2} \end{bmatrix} \begin{bmatrix} \hat{a}(\tilde{\Delta}) \\ \hat{a}^\dagger(-\tilde{\Delta}) \end{bmatrix}, \quad (3.123)$$

where ϕ is the phase of the pump-induced field $\alpha = |\alpha|e^{i\phi}$ inside the cavity. If we choose to regard $\tilde{\Delta}$ as the signal frequency, $-\tilde{\Delta}$ is referred to as the idler frequency. Inverting Eq. (3.123) and using the boundary condition, Eq. (3.100), we end up with our final Fourier domain relationship between the incoming and outgoing signal fields:

$$\hat{a}_{\text{out}}(\tilde{\Delta}) = g_s(\tilde{\Delta})\hat{a}_{\text{in}}(\tilde{\Delta}) + g_i(\tilde{\Delta})\hat{a}_{\text{in}}^\dagger(-\tilde{\Delta}), \quad (3.124)$$

where

$$g_s(\tilde{\Delta}) = 1 - \frac{i(\delta - 2\xi\tilde{n} - \tilde{\Delta}) + 1/2}{(i\tilde{\Delta} - \lambda_-)(i\tilde{\Delta} - \lambda_+)} \quad (3.125)$$

$$g_i(\tilde{\Delta}) = \frac{\xi\tilde{n}e^{2i\phi}}{(i\tilde{\Delta} - \lambda_-)(i\tilde{\Delta} - \lambda_+)} \quad (3.126)$$

and

$$\lambda_{\pm} = \frac{1}{2} \pm \sqrt{(\xi\tilde{n})^2 - (\delta - 2\xi\tilde{n})^2}. \quad (3.127)$$

The frequency-dependent signal and idler amplitude gains $g_s(\tilde{\Delta})$ and $g_i(\tilde{\Delta})$ relate the output field to the input field at $\tilde{\Delta}$ and to its conjugate at $-\tilde{\Delta}$. The signal and idler power gains⁶¹ are

$$G_s(\tilde{\Delta}) = |g_s(\tilde{\Delta})|^2 \quad (3.128)$$

$$G_i(\tilde{\Delta}) = |g_i(\tilde{\Delta})|^2 = G_s(\tilde{\Delta}) - 1 \quad (3.129)$$

The Kerr nonlinearity has mapped the behavior of tones input at both $\pm\tilde{\Delta}$, equally spaced above and below the pump, to a single output frequency. In other words, the JPA's output at a given signal frequency ω cares not only about the signal input at ω , but about the input at the idler frequency, $2\omega_p - \omega$. Since the reverse is true for the output at $2\omega_p - \omega$, there is a spectral ambiguity in the output:⁶² a signal at a given frequency could have originated at either relevant input frequency.

3.4.5 Gain-bandwidth product

Another important property of the JPA is its preservation of the amplitude gain-bandwidth product (GBP). The signal and idler gains are both symmetric about the pump frequency,

$$G_s(\tilde{\Delta}) = G_s(-\tilde{\Delta}) \quad (3.130)$$

$$G_i(\tilde{\Delta}) = G_i(-\tilde{\Delta}) \quad (3.131)$$

⁶¹ Both the signal and idler gains are phase-insensitive, though not phase-preserving, as the idler gain is phase-conjugating. The signal gain is also referred to as the “direct gain.” It amplifies both quadratures of the microwave field equally, so where that property is important to distinguish from single-quadrature gain, (which is a linear combination of signal and idler gains, Sec. 3.4.6), I will refer to it as the “two-quadrature gain.” Additionally, though the signal/direct/two-quadrature gain is a frequency-dependent quantity per Eq. (3.128), in some contexts it will be convenient to refer to the maximum value $G_s(0)$ as simply the signal/direct/two-quadrature gain.

⁶² The -1 in Eq. (3.129) means the contributions are not strictly identical, but for typical $G_s \gtrsim 100$, the difference is quite negligible to amplification. For squeezing, however, it is important: without it, squeezing would be infinite at any gain.

with full-width at half-max JPA bandwidth

$$B = 2\lambda_+. \quad (3.132)$$

Near the critical point, Eqs. (3.117)–(3.119),

$$\sqrt{G_s(0)}B \approx \frac{2}{\sqrt{3}}\gamma_e. \quad (3.133)$$

The quality factor of the resonant circuit in Fig. 3.6a therefore sets the bandwidth (for a given gain) of the Kerr resonator in Fig. 3.6c, independent of nonlinearity.

3.4.6 Single-quadrature amplification, deamplification, and squeezing

We can readily see that Eq. (3.124) gives rise to squeezing. Identifying

$$g_s(\tilde{\Delta}) = \cosh r \quad (3.134)$$

$$g_i(\tilde{\Delta}) = e^{i\theta} \sinh r, \quad (3.135)$$

Eq. (3.124) can be recast as

$$\hat{a}_{\text{out}}(\tilde{\Delta}) = \hat{S}^\dagger(re^{i\theta})\hat{a}_{\text{in}}(\tilde{\Delta})\hat{S}(re^{i\theta}) = \hat{a}_{\text{in}}(\tilde{\Delta}) \cosh r - \hat{a}_{\text{in}}^\dagger(-\tilde{\Delta})e^{i\theta} \sinh r, \quad (3.136)$$

whereupon we see $\hat{S}(re^{i\theta})$ to be the unitary squeezing operator [193, 194]. If we define the Fourier domain quadrature operators for the signal as

$$\hat{X}(\tilde{\Delta}) = \frac{1}{\sqrt{2}} \left(\hat{a}(\tilde{\Delta}) + \hat{a}^\dagger(-\tilde{\Delta}) \right) \quad (3.137)$$

$$\hat{Y}(\tilde{\Delta}) = \frac{1}{\sqrt{2}i} \left(\hat{a}(\tilde{\Delta}) - \hat{a}^\dagger(-\tilde{\Delta}) \right), \quad (3.138)$$

and set $\theta = 0$, the squeezing operator reduces the variance of $\hat{X}(\tilde{\Delta})$ for an input vacuum state $|0\rangle$ below its vacuum level by e^{-2r} and amplifies that of $\hat{Y}(\tilde{\Delta})$ by e^{2r} . For general θ , the same applies to a pair of generalized (i.e. rotated) quadratures.

The two-quadrature gain G_s and the single-quadrature gain or deamplification⁶³ G_{1Q} parametrized by θ are related by the equation

$$G_{1Q}(G_s, \theta) = (2G_s - 1) + 2\sqrt{G_s(G_s - 1)} \cos(2\theta), \quad (3.139)$$

⁶³ Deamplification is a more general term than squeezing for what the JPA does to the quadrature out of phase with the pump. When the input state has its variance deamplified below vacuum, it is said to be squeezed.




Figure 3.9: (a) Single-quadrature gain G_{1Q} plotted as a function of pump-probe phase difference for a two-quadrature gain $G_s = 20$. (b) Closeup of (a) for $G_{1Q} < 1$, where deamplification occurs.

plotted in Fig. 3.9a. Figure 3.9b shows the region below $G_{1Q} = 1$, where deamplification occurs. Maximum single-quadrature amplification (deamplification) is achieved at $\theta = \{0, \pi\}$ ($\{\pi/2, 3\pi/2\}$), where in the high signal gain ($G_s \gg 1$) limit

$$G_{1Q}(G_s \rightarrow \infty, 0) \rightarrow 4G_s \quad (3.140)$$

$$G_{1Q}\left(G_s \rightarrow \infty, \frac{\pi}{2}\right) \rightarrow \frac{1}{4G_s}. \quad (3.141)$$

Because of the phase-dependent behavior of single-quadrature gain, this mode of operation is also referred to as phase-sensitive, and conversely two-quadrature gain occurs with phase-insensitive operation.

When performing phase-insensitive amplification, the JPA must add at least one half quantum of noise spectral density distributed equally over both quadratures, the minimum allowed

by quantum mechanics for two-quadrature measurement. [128, 199, 200]. This noise traces to the vacuum fluctuations of the idler mode,⁶⁴ now mapped to the signal output. In phase-sensitive operation, the amplifier in principle does not have to add any noise along the amplified quadrature, making the JPA an excellent tool for measuring feeble signals.

This concludes our discussion of current-pumped JPA amplification, deamplification, and squeezing. In Sec. 3.4.7, we will see that flux pumping can give rise to the same phenomena, without explicit reliance on the Kerr nonlinearity.

3.4.7 Derivation of flux-pumped JPA gain

Relative to current-pumped JPAs, flux-pumped JPAs perform a more idealized operation wherein the inductance of the SQUID(s) can be thought of as simply varying at the pump frequency. When that frequency is twice the resonant frequency, parametric amplification is realized [141, 182], as in the current-pumped case treated in Sec. 3.4.4. This section closely follows the derivation in Appx. A of Ref. [201], with the goal of obtaining the signal and idler gains equivalent to Eqs. (3.128) and (3.129).

We begin with the Hamiltonian of Eq. (3.86), which we recast in terms of the TKC resonant frequency ω_{TKC} , Eq. (3.73), as

$$\mathcal{H}_0(\Phi_a, Q_\Delta) = \frac{Q_\Delta^2}{2C_\Sigma} + \frac{C_\Sigma}{2} \omega_{\text{TKC}}^2 \Phi_a^2. \quad (3.142)$$

We will modulate the resonance as

$$\omega_{\text{TKC}} \rightarrow \omega_{\text{TKC}} [1 + \Lambda \cos(\omega_p t)], \quad (3.143)$$

where $\Lambda \ll 1$ and ω_p are the pump amplitude and frequency, respectively.⁶⁵ Neglecting the Λ^2 term, Eq. (3.143) modifies the Hamiltonian as

$$\mathcal{H}_0(\Phi_a, Q_\Delta) = \frac{Q_\Delta^2}{2C_\Sigma} + \frac{C_\Sigma}{2} \omega_{\text{TKC}}^2 [1 + \epsilon \cos(\omega_p t)] \Phi_a^2, \quad (3.144)$$

⁶⁴ That is, it corresponds to the $+1/2$ in the idler mode's harmonic oscillator Hamiltonian (Footnote 45).

⁶⁵ In practice, Eq. (3.143) is the dominant contribution from varying the array inductance, Eq (3.87), at ω_p . In particular, a time-varying Φ_{ext} gives rise to a time-varying Φ_{tot} via Eq. (3.38), which varies the SQUID critical currents, Eq. (3.27), in order to vary the array inductance and with it ω_{TKC} .

where

$$\epsilon = 2\Lambda. \quad (3.145)$$

The corresponding equation of motion of the resonator field operator \hat{A} , where we have promoted the canonical variables to operators as in Sec. 3.3.4, is given by

$$\frac{d\hat{A}}{dt} = \left(-i\omega_{\text{TKC}} - \frac{\gamma_e}{2}\right)\hat{A} - 2i\omega_{\text{TKC}}\epsilon \cos(\omega_p t)(\hat{A} + \hat{A}^\dagger) + \sqrt{\gamma_e}\hat{A}_{\text{in}}, \quad (3.146)$$

where we have ignored the ground state energy term and used the operator-equivalents of Eqs. (3.90) and (3.91) for the field operator and its conjugate. We rewrite Eq. (3.146) as

$$\frac{d\hat{A}}{dt} + W\hat{A} + V\hat{A}^\dagger + \hat{F}, \quad (3.147)$$

where

$$W(t) = i\omega_{\text{TKC}} + \frac{\gamma_e}{2} + 2i\omega_{\text{TKC}}\epsilon \cos(\omega_p t) \quad (3.148)$$

$$V(t) = 2i\omega_{\text{TKC}}\epsilon \cos(\omega_p t) \quad (3.149)$$

$$\hat{F}(t) = \sqrt{\gamma_e}\hat{A}_{\text{in}}(t) \quad (3.150)$$

show the time dependencies of each term explicitly. From Eq. (3.147), we obtain the three equations

$$\frac{d\hat{A}^\dagger}{dt} = -W^*\hat{A}^\dagger - V^*\hat{A} + \hat{F}^\dagger \quad (3.151)$$

$$\hat{A}^\dagger = \frac{1}{V} \left(\hat{F} - \frac{d\hat{A}}{dt} - W\hat{A} \right) \quad (3.152)$$

$$\frac{d^2\hat{A}}{dt^2} + \frac{dW}{dt}\hat{A} + W\frac{d\hat{A}}{dt} + \frac{dV}{dt}\hat{A}^\dagger + V\frac{d\hat{A}^\dagger}{dt} = \frac{d\hat{F}}{dt}. \quad (3.153)$$

Together Eqs. (3.151)–(3.153) yield

$$\begin{aligned} \frac{d^2\hat{A}}{dt^2} + \left(W + W^* - \frac{1}{V}\frac{dV}{dt}\right)\frac{d\hat{A}}{dt} + \left(\frac{dW}{dt} - \frac{W}{V}\frac{dV}{dt} + |W|^2 - |V|^2\right)\hat{A} \\ = \frac{d\hat{F}}{dt} + \left(W^* - \frac{1}{V}\frac{dV}{dt}\right)\hat{F} - V\hat{F}^\dagger. \end{aligned} \quad (3.154)$$

We now specify to the case of a classical resonator field and signal input

$$\hat{A} \rightarrow \alpha \quad (3.155)$$

$$\hat{A}_{\text{in}} \rightarrow Ee^{i\omega_s t}, \quad (3.156)$$

which will be sufficient to derive the amplification properties of the flux-pumped JPAs. In Eq. (3.156), E and ω_s are the signal amplitude and frequency, respectively. Performing an RWA by letting terms rotating rapidly relative to $\omega_p \approx 2\omega_{\text{TKC}}$ and $\omega_s \approx \omega_{\text{TKC}}$ average away, Eq. (3.154) becomes

$$\begin{aligned} & \frac{d^2\alpha}{dt^2} + [\kappa + i\omega_p] \frac{d\alpha}{dt} + \left[\frac{\gamma_e^2}{4} + \omega_{\text{TKC}}^2 (1 - \epsilon^2 - \frac{\omega_p}{\omega_{\text{TKC}}}) + i\omega_p \frac{\gamma_e}{2} \right] \alpha \\ &= \sqrt{\gamma_e} E \left[\frac{\gamma_e}{2} + i(\omega_p - \omega_s - \omega_{\text{TKC}}) \right] e^{-i\omega_s t} - i\omega_{\text{TKC}} \epsilon \sqrt{\gamma_e} E^* e^{-i(\omega_p - \omega_s)t}. \end{aligned} \quad (3.157)$$

To solve Eq. (3.157), we write down the homogeneous solution, obtained by setting the right-hand side to zero, and the particular solution. The homogeneous solution takes the form

$$\alpha_h = C_+ e^{\eta_+ t} + C_- e^{\eta_- t}, \quad (3.158)$$

where

$$\eta_{\pm} = -\frac{\gamma_e + i\omega_p}{2} \pm \omega_{\text{TKC}} \sqrt{\epsilon^2 - \left(\frac{\omega_p}{2\omega_{\text{TKC}}} - 1 \right)^2}. \quad (3.159)$$

The particular solution is

$$\alpha_p = -y_+ \int \frac{f y_-}{D} dt + y_- \int \frac{f y_+}{D} dt, \quad (3.160)$$

where

$$y_{\pm} = e^{\eta_{\pm} t} \quad (3.161)$$

$$D = (\eta_- - \eta_+) e^{(\eta_- + \eta_+) t} \quad (3.162)$$

$$f = \gamma_e E \left[\frac{\gamma_e}{2} + i(\omega_p - \omega_s - \omega_{\text{TKC}}) \right] e^{-i\omega_s t} - i\omega_{\text{TKC}} \epsilon \sqrt{\gamma_e} E^* e^{-i(\omega_p - \omega_s)t}. \quad (3.163)$$

Hence

$$\alpha_p = A_s e^{-i\omega_s t} + A_i e^{-i(\omega_p - \omega_s)t}, \quad (3.164)$$

where

$$A_s = \frac{\sqrt{\gamma_e} E \left[\frac{\gamma_e}{2} + i(\omega_p - \omega_s - \omega_{\text{TKC}}) \right]}{\left[\frac{\gamma_e}{2} + i(\omega_{\text{TKC}} - \omega_s) \right] \left[\frac{\gamma_e}{2} + i(\omega_p - \omega_s - \omega_{\text{TKC}}) \right] - \epsilon^2 \omega_{\text{TKC}}^2} \quad (3.165)$$

$$A_i = \frac{i\omega_{\text{TKC}} \epsilon \sqrt{\gamma_e} E^*}{\left[\frac{\gamma_e}{2} + i(\omega_{\text{TKC}} - \omega_s) \right] \left[\frac{\gamma_e}{2} + i(\omega_p - \omega_s - \omega_{\text{TKC}}) \right] - \epsilon^2 \omega_{\text{TKC}}^2}. \quad (3.166)$$

The full solution to Eq. (3.157) is

$$\alpha = \alpha_h + \alpha_p. \quad (3.167)$$

For the homogeneous solution α_h to die out, the real parts of both η_+ and η_- need to be negative. While $\text{Re}(\eta_-) < 0$ always, $\text{Re}(\eta_+) < 0$ occurs when

$$\frac{\gamma_e}{2} > \omega_{\text{TKC}} \sqrt{\epsilon^2 - \left(\frac{\omega_p}{2\omega_{\text{TKC}}} - 1 \right)^2}. \quad (3.168)$$

We define the critical value

$$\epsilon_c = \frac{\gamma_e}{2\omega_{\text{TKC}}} = \frac{1}{2Q_e}, \quad (3.169)$$

where the last equality uses Eq. (3.77). For pumping at twice the resonant frequency, $\omega_p = 2\omega_{\text{TKC}}$, ϵ_c is seen via Eq. (3.145) to be twice the modulation amplitude of the resonance frequency where the flux-pumped JPA response diverges. This corresponds to modulating the TKC resonance, Eq. (3.143), by a fractional amount

$$\Lambda \approx \frac{1}{4Q_e} \quad (3.170)$$

in order to achieve high gain.

Operating below the critical point,⁶⁶ the steady state solution to Eq. (3.157) is α_p , Eq. (3.164). Using our boundary condition, Eq. (3.100), the classical output signal is given by

$$A_{\text{out}} = (E - \sqrt{\gamma_e} A_s) e^{-i\omega_s t} - \sqrt{\gamma_e} A_i e^{-(\omega_p - \omega_s)t}. \quad (3.171)$$

For the non-degenerate case where $\omega_p \neq 2\omega_s$, Eq. (3.171) corresponds to signal and idler power gains

$$G_s = \left| 1 - \frac{\sqrt{\gamma_e} A_s}{E} \right|^2 = \left| 1 - \gamma_e \frac{\left[\frac{\gamma_e}{2} + i(\omega_p - \omega_s - \omega_{\text{TKC}}) \right]}{\left[\frac{\gamma_e}{2} + i(\omega_{\text{TKC}} - \omega_s) \right] \left[\frac{\gamma_e}{2} + i(\omega_p - \omega_s - \omega_{\text{TKC}}) \right] - \epsilon^2 \omega_{\text{TKC}}^2} \right|^2 \quad (3.172)$$

$$G_i = \left| \frac{\sqrt{\gamma_e} A_i}{E^*} \right|^2 = \left| \frac{\gamma_e \omega_{\text{TKC}} \epsilon}{\left[\frac{\gamma_e}{2} + i(\omega_{\text{TKC}} - \omega_s) \right] \left[\frac{\gamma_e}{2} + i(\omega_p - \omega_s - \omega_{\text{TKC}}) \right] - \epsilon^2 \omega_{\text{TKC}}^2} \right|^2. \quad (3.173)$$

These gains are the analogues of Eqs. (3.128) and (3.129) for the current-pumped G_s and G_i in Sec. 3.4.4. As in Secs. 3.4.5 and 3.4.6, the flux-pumped JPA has a GBP set by its external coupling rate γ_e , and can be used to amplify or deamplify one quadrature of an input signal. In

⁶⁶ For a derivation of the bifurcation properties of the flux-pumped JPA operated above its critical point, see the Supplemental Material of Ref. [170].

Chs. 5 and 8, we will see them used experimentally to squeeze and amplify vacuum in the service of speeding up dark matter axion detection efforts.

Having now derived the key amplification properties of both current- and flux-pumped JPAs, we will return in Sec. 3.5 to the current-pumped case in order to examine the role of stray inductance and the use of SQUID arrays (Fig. 3.6a) in the TKC design. The ensuing discussion will inform our design of current- as well as flux-pumped JPAs in Secs. 3.6 and 3.7, respectively.

3.5 Validity of the Duffing approximation

There is cause to question the validity of the Duffing approximation [195–197] for describing the physics of array-based, current-pumped JPAs (Henceforth, array-based JPAs).⁶⁷ Recall that in the Duffing (or Kerr) approximation shortened the full Hamiltonian, Eq. (3.84), to that of a linear LC resonator plus the first nonlinear correction due to the SQUID array, Eq. (3.85). That assumption was predicated on small excitations (in terms of flux relative to a flux-quantum, or equivalently current relative to critical current) of the SQUIDs. In their 2007 paper, Manucharyan *et al.* [181] derive the requirements on circuits used to embed Josephson junctions in order to create parametric amplifiers. Their central requirement for the validity of the Duffing approximation is violated by our array-based JPAs. At stake is not whether array-based JPAs really work: theoretical and experimental evidence indicate strongly that they do [148, 163, 188, 198]. Rather, the question is whether they work in spite of violating the Duffing approximation, or subject to it. This section presents the requirement of Ref. [181], which applies directly to single-junction⁶⁸ JPAs in the presence of geometrical inductances, extends it to array-based JPAs, and finally shows that it is incorrect, even for single-junction JPAs. In its place, we will be left with a modified version of the violated condition of Ref. [181] which our array-based JPAs are seen not to violate.

⁶⁷ The key point here is we are assuming the oscillator is operated near the critical point, as JPAs typically are (Sec. 3.4.4), in order to inquire if the Duffing approximation adequately describes them in that limit. The Duffing approximation is of course perfectly valid for low-amplitude drives.

⁶⁸ The difference between series arrays of junctions and series arrays of SQUIDs is immaterial to this section. For both simplicity and proximity to Ref. [181], this section works with junctions and arrays thereof.

3.5.1 The Q_p condition for a single-junction JPA

Manucharyan *et al.* [181] analyze a single Josephson junction embedded in an arbitrary linear microwave network, shown in a parallel and series configuration in Fig. 3.10. For a junction with critical current I_c , the constitutive equation, Eq. (3.3), expands as

$$I = I_c \sin \frac{\Phi}{\phi_0} = I_c \frac{\Phi}{\phi_0} \left[1 - \frac{1}{6} \left(\frac{\Phi}{\phi_0} \right)^2 + \dots \right], \quad (3.174)$$

where the first nonlinear term is the Duffing term.




Figure 3.10: (a) The Josephson junction is represented up through its first nonlinear term as an ideal inductor in parallel with a parallel nonlinear element (PNL) governed by Eq. (3.175). (b) The embedding environment for the PNL is modeled by an admittance $Y(\omega)$ with drive and noise current sources represented by their Norton equivalents. (c), (d), and Eq. (3.176) are the series nonlinear element (SNL)/Thévenin equivalent versions of (a), (b), and Eq. (3.175), respectively. Panels (b) and (d) reproduced from Ref. [181].

The relevant procedure is as follows: first, a decision is made regarding the use of the parallel nonlinear element (PNL) versus the series nonlinear element (SNL), shown in Fig. 3.10a and c, respectively. Using the PNL (SNL) requires modeling the drive as its Norton (Thévenin) equivalent circuit, as in Fig. 3.10b(d). The linear portion of Josephson inductance L_{J0} , Eq. (3.10), is folded into the equivalent admittance $Y(\omega)$ (impedance $Z(\omega)$). The cubic constitutive equations for the

PNL and SNL follow directly from Eq. (3.174):

$$I_{\text{PNL}}(t) = -\frac{1}{6L_{J0}\phi_0^2}\Phi_{\text{PNL}}^3(t) \quad (3.175)$$

$$\Phi_{\text{SNL}}(t) = \frac{L_{J0}^3}{6\phi_0^2}I_{\text{SNL}}^3(t). \quad (3.176)$$

With a model chosen, the next step is to define a participation ratio⁶⁹ $p \in [0, 1]$, which characterizes the amount of Josephson inductance relative to total linear inductance in the embedding circuit (in which L_{J0} is now included). In Fig. 3.6a, additional linear inductance is represented by L_G . Reference [181] gives simple examples for the cases of a parallel and a series RLC circuit, shown in Fig. 3.11a and b, respectively. Calling the linear inductance L_p for the parallel case or L_s for the series case, $p_p = L_p/L_{J0}$ for the parallel case and $p_s = L_{J0}/L_s$ for the series case. The other relevant parameter to the analysis is the embedding circuit quality factor Q .




Figure 3.11: (a) Current-driven parallel RLC circuit with a PNL. (b) Voltage-driven series RLC circuit with an SNL. (a) and (b) correspond to simple cases of Fig. 3.10b and d, respectively. Figure reproduced from Ref. [181].

Reference [181] analyzes the simple and general cases, and concludes that

$$Qp \gg 1, \quad (3.177)$$

⁶⁹ The subscript-less p is used to denote the participation ratio for statements that refer to both the series and parallel cases.

is the necessary condition to keep the maximum current through the junction well below I_c , and hence to satisfy the Duffing approximation.

3.5.2 Significance of the participation ratio

The embedding circuit machinery relies upon the notion that the linear portion of the Josephson inductance can be wrapped up into a linear inductance apart from the junction itself without any change to the physics. The Josephson inductance as it appears in the participation ratio is there not merely as a quantification of the linear inductance that the junction carries with it, but as a proxy for the nonlinear inductance that is always associated with that particular linear inductance. This, then, is the significance of the ratio of Josephson linear inductance-to-non-Josephson linear inductance: the Josephson inductance always carries with it, *in fixed ratio*, a nonlinear inductance, while any stray linear inductance does not.

3.5.3 From junction to array

The ratio of nonlinear-to-linear inductance changes when the junction is swapped out for a series junction array. In order to faithfully map the participation ratio from a single junction to an equivalent array (meaning one with the same linear inductance, which will therefore give the same bare resonance), we must first determine the composition of that equivalent array. Then we will use the fact that we can change what linear inductance we choose to call Josephson or stray⁷⁰ in order to restore the nonlinear-to-linear inductance ratio of the single-junction case, whence the participation ratio is immediately obtained. We will follow this procedure for two cases: the series *RLC* case (which generalizes to the general Thévenin/impedance case, Fig. 3.10b) and the parallel *RLC* case (likewise the general Norton/admittance case, Fig. 3.10d). The result is the same in both cases: the participation ratio decreases quadratically with N .

The constitutive equation of a series array of N junctions (each of critical current I'_c and linear Josephson inductance $L'_J = \phi_0/I'_c$) differs fundamentally from that of a single junction or

⁷⁰ The non-Josephson embedding inductance need not be undesirable, as the word “stray” might imply.

parallel array, as N infiltrates the sine:⁷¹

$$I = I'_c \sin \frac{\Phi_a}{N\phi_0} = I'_c \frac{\Phi_a}{N\phi_0} \left[1 - \frac{1}{6} \left(\frac{\Phi_a}{N\phi_0} \right)^2 + \dots \right], \quad (3.178)$$

where I is the current through and Φ_a the flux across the array. The nonlinearity has been reduced relative to the linear Josephson inductance by a factor of N^2 .

3.5.4 Series RLC and general Thévenin case

In the case of a series RLC embedding circuit, we are concerned with the stray inductance $L_{s,\text{stray}}$ together with the linear and nonlinear inductances of the array.⁷² We begin by solving Eq. (3.178) for Φ_a to third order in I , obtaining

$$\Phi_a = N\phi_0 \arcsin \frac{I}{I'_c} = NL'_{J0} I + \frac{1}{6} NL'_{J0} \left(\frac{I}{I'_c} \right)^2 I + \dots \quad (3.179)$$

Dividing by I gives the linear (L_{a0}) and first nonlinear ($L_{s,\text{nl}}$) inductance terms:

$$L_{a0} = NL'_{J0} \quad (3.180)$$

$$L_{s,\text{nl}} = \frac{1}{6} NL'_{J0} \left(\frac{I}{I'_c} \right)^2. \quad (3.181)$$

As discussed in Sec. 3.5.2, in order for the participation ratio to be a meaningful proxy for the nonlinear inductance, we need to ensure that the linear and nonlinear array inductances maintain a fixed ratio r_s :

$$r_s = \frac{L_{s,\text{nl}}}{L_{a0}} = \frac{1}{6} \left(\frac{I}{I'_c} \right)^2. \quad (3.182)$$

For the array and junction to be equivalent, they must have the same linear inductance. We therefore require

$$L_{a0} = NL'_{J0} = L_{J0}. \quad (3.183)$$

⁷¹ In the case of a parallel array, such as the simple case of a SQUID, adding junctions in parallel increases the critical current, Eq. (3.32), something that in principle can be done with a different physical junction, Eq. (3.1). Adding junctions in series, conversely, is like changing the flux quantum: the constitutive equation of a series array cannot be achieved with any single junction; as we will see, however, it can be mimicked by a single junction plus some linear inductance.

⁷² If we take there to be an inductive term, which we label $L_{s,\text{stray}}$, somewhere in the general Thévenin case, Fig. 3.10d, the following analysis will apply equally to the more general case.

Since ϕ_0 is a constant of nature, the only way to satisfy Eq. (3.183) is with $I'_c = NI_c$. Written in terms of I_c , Eq. (3.182) becomes

$$r_s = \frac{1}{6N^2} \left(\frac{I}{I_c} \right)^2. \quad (3.184)$$

We see that I'_c 's dependence on N , brought about to ensure equal array and single-junction linear inductances, implies an N -dependent inductance ratio. The Josephson inductance is now no longer the proxy for nonlinear inductance it was in the single-junction case, causing the participation ratio to no longer carry the same meaning. We will restore the constancy of this ratio for the series and then the parallel embedding circuit case with some algebra incorporating stray linear inductances.




Figure 3.12: Series inductor algebra showing the equivalence between stray and array linear inductances. The total series inductance of an N -Josephson junction array and a stray series inductance $L_{s,\text{stray}}$ (top left) is represented to first nonlinear order by writing the array's inductance in terms of its linear and first nonlinear contributions (top right). An equivalent configuration (bottom right), consisting of the same total linear inductance and nonlinear-to-linear array inductance ratio r_s , is achieved when going from array to single-junction by increasing the value of the stray inductance.

To do so, we start by equating the array of junctions with a linear and nonlinear inductor, as shown in Fig. 3.12. Then we use the physical indistinguishability of the array's linear inductance and the stray linear inductance to move some of the former over to the latter such that the total linear inductance is preserved and the ratio of nonlinear-to-Josephson inductance is modified to match the $N = 1$ case. The two effective linear inductances become:

$$L_{a0} \rightarrow L_{s,a}^{\text{eff}} = \frac{1}{N^2} L_{a0} \quad (3.185)$$

$$L_{s,\text{stray}} \rightarrow L_{s,\text{stray}}^{\text{eff}} = L_{s,\text{stray}} + \frac{N^2 - 1}{N^2} L_{a0}. \quad (3.186)$$

The sum of linear inductances is preserved, and the ratio

$$r_s \rightarrow r_s^{\text{eff}} = \frac{L_{s,\text{nl}}}{L_{s,a0}^{\text{eff}}} = \frac{1}{6} \left(\frac{I}{I_c} \right)^2 \quad (3.187)$$

is N -independent, as desired.

The array inductance is *now* a proxy for the nonlinear inductance in just the manner that the single-junction inductance is in the simpler case. Hence, Ref. [181]’s conclusions pertaining to p will hold exactly if we redefine p_s :

$$p_s = \frac{L_{J0}}{L_s} = \frac{L_{J0}}{L_{J0} + L_{s,\text{stray}}} \rightarrow p_s^{\text{eff}} = \frac{L_{s,a0}^{\text{eff}}}{L_s^{\text{eff}}} = \frac{L_{s,a0}^{\text{eff}}}{L_{s,a0}^{\text{eff}} + L_{s,\text{stray}}^{\text{eff}}} = \frac{p_s}{N^2}, \quad (3.188)$$

where L_s and $L_{s,\text{eff}}$ are the total linear series inductance and are equal to one another. We conclude that replacing a single junction with an equally linearly inductive array of N junctions reduces the participation ratio by a factor of N^2 in the series embedding case.

3.5.5 Parallel RLC and general Norton case




Figure 3.13: Parallel inductor algebra showing the equivalence between stray and array linear inductances. The parallel case reasons analogously to the series case of Fig. 3.12: the array (left) is expressed as a linear inductor L_{a0} and a nonlinear inductor $L_{p,\text{nl}}$ (center). The linear inductance is increased to restore the ratio r_p to its $N = 1$ value when going from array to single-junction. Doing so necessitates altering the stray inductance (right).

The reasoning for the parallel case, shown in Fig. 3.13, is analogous to the series case. To first nonlinear order, the array is represented by a linear inductance $L_{a0} = N L'_{J0}$ (where, as before,

$NL'_{J0} = L_{J0} \rightarrow I'_c = NI_c$) in parallel with a nonlinear inductance $L_{p,nl} = r_p L_{a0}$, with the ratio r_p given by

$$r_p = 6N^2 \left(\frac{\phi_0}{\Phi_a} \right)^2. \quad (3.189)$$

As such, L_{a0} needs to be increased by N^2 :

$$L_{a0} \rightarrow L_{p,a0}^{\text{eff}} = N^2 L_{a0}, \quad (3.190)$$

in order to restore the $N = 1$ ratio and give the participation ratio its original meaning of Sec. 3.5.2.

To preserve the parallel linear inductance, we also require

$$L_{p,\text{stray}} \rightarrow L_{p,\text{stray}}^{\text{eff}} = \frac{N^2 L_{a0} L_{p,\text{stray}}}{(N^2 - 1)L_{p,\text{stray}} + N^2 L_{a0}}, \quad (3.191)$$

yielding

$$p_p = \frac{L_p}{L_{J0}} = \frac{\left(\frac{1}{L_{a0}} + \frac{1}{L_{p,\text{stray}}} \right)^{-1}}{L_{p,a0}} \rightarrow p_p^{\text{eff}} = \frac{L_p^{\text{eff}}}{L_{a0}^{\text{eff}}} = \frac{\left(\frac{1}{L_{p,a0}^{\text{eff}}} + \frac{1}{L_{p,\text{stray}}^{\text{eff}}} \right)^{-1}}{L_{p,a0}^{\text{eff}}} = \frac{p_p}{N^2}, \quad (3.192)$$

the same result as in the series embedding case.

3.5.6 The Qp condition for an array-based JPA

For both the series and parallel embedding circuits, Eqs. (3.188) and (3.192) indicate that Eq. (3.177) becomes

$$Qp \gg N^2, \quad (3.193)$$

where p is the original participation ratio defined between the linear Josephson (more generally, array) and total series or parallel linear inductance. Typical JPAs with $N \sim 10$ junctions (or SQUIDS) and $Q \sim 50$ fail to meet this condition, yet remain functional.

We will see that array-based JPAs work because the participation ratio in fact has no place in what would otherwise be the Qp -condition. In concert with the findings above, this implies that stray inductance can, like SQUID arrays, be used to increase power handling [148], but potentially without the adverse consequences of nonuniform junction or SQUID loop fabrication [173].

3.5.7 Independence of the validity of the Duffing approximation from the nonlinearity

The Duffing approximation analysis in Ref. [181] maps back to the equation for a damped, driven Duffing oscillator,

$$\ddot{X} + \frac{\omega_0}{Q}\dot{X} + \omega_0^2 X(1 - \nu X^2) = A \cos \omega_d t, \quad (3.194)$$

with X the coordinate, ω_0 the bare resonance frequency, ν the strength of the nonlinearity, and the right-hand side a drive tone of amplitude A and frequency ω_d . For an electrical system, the participation ratio p maps directly onto the nonlinearity. For example, in the case of the parallel *RLC* circuit with a single junction, ν corresponds to $p/6$.

Let us consider the effect of changing ν (equivalently p) on the solution of Eq. (3.194) and on that of the fully nonlinear equation from which it comes,

$$\ddot{X} + \frac{\omega_0}{Q}\dot{X} + \frac{\omega_0^2}{\sqrt{6\nu}} \sin(X\sqrt{6\nu}) = A \cos \omega_d t. \quad (3.195)$$

The sine in Eq. (3.195) can be expanded, and an arbitrary number n of nonlinear terms kept:

$$\ddot{X} + \frac{\omega_0}{Q}\dot{X} + \omega_0^2 X \sum_{m=0}^n \frac{(-6)^m}{(2m+1)!} (X^2\nu)^m = A \cos \omega_d t. \quad (3.196)$$

Equation (3.194) corresponds to $n = 1$, and Eq. (3.195) to $n = \infty$.

Inspection reveals that, if Eq. (3.196) for arbitrary n and ν has solution $X_n^\nu(t)$, then the equation for n and ν' will be solved by $\sqrt{\nu/\nu'} X_n^\nu(t)$, provided that the drive amplitude is likewise adjusted, $A \rightarrow \sqrt{\nu/\nu'} A$. This corresponds to driving the JPA at a fixed percentage of its critical power.⁷³

Therefore, the ratio of any two solutions for arbitrary n_1 and n_2 and a given nonlinearity ν ,

$$Z_{n_1, n_2}^\nu(t) \equiv \frac{X_{n_1}^\nu(t)}{X_{n_2}^\nu(t)}, \quad (3.197)$$

remains unaffected if the nonlinearity is changed to ν' . The significance of this statement stems from the fact that $Z_{1, \infty}^\nu(t)$ directly quantifies the divergence of the Duffing approximation from the

⁷³ This can also be seen in Eq. (3.114), wherein the Kerr constant only enters alongside the drive strength [148].

full solution: in other words, it measures how valid the Duffing approximation is. The fact that $Z_{1,\infty}^\nu(t)$ is independent of ν tells us that the nonlinearity does not affect the validity of the Duffing approximation for JPAs.⁷⁴ So the Duffing approximation may be however bad or good it is, but its correctness does not depend on changes to the nonlinearity.

3.5.8 Modified condition for the Duffing approximation

Since Q does not ultimately affect the dynamics of the system, Eq. (3.195) or Eq. (3.196), in the same fashion as ν does (its situation alongside p in the Qp product, Eq. (3.177) or (3.193), being therefore somewhat incidental in this sense), there is no reason that the Duffing approximation's validity for JPA operation should be independent of Q . Numerical tests which confirm the prediction of independence of Duffing validity from ν reveal that the Duffing approximation does break down as Q approaches unity from above. The phase response of a reflected tone b_{out} is shown near the critical point in Fig. 3.14 for different orders of approximation $n = \{1, 2, \infty\}$ and values of Q off of a lossless, array-based JPA. The array of $N = 20$ junctions easily violates the modified Qp condition, Eq. (3.193), for all values of Q shown. Nevertheless, the agreement between the Duffing system ($n = 1$) and the full sinusoidal solution ($n = \infty$) is quite good at high Q . This confirms that the participation ratio plays no role in determining the validity of the Duffing approximation for JPA operation, but that Q still does. Thus,

$$Q \gg 1, \tag{3.198}$$

independent of geometrical inductance or number of junctions, is the relevant predictor of Duffing approximation validity for JPA operation. The presence of stray inductance or of Josephson junction arrays have no bearing.

This concludes our detour into the theory of the Duffing oscillator. The independence of the validity of the Duffing approximation for JPAs from the degree of nonlinearity in the system explains why no divergence from the Duffing results has been noted for array-based JPAs in the literature.

⁷⁴ More generally, it does not affect the agreement of any two levels of approximation, parametrized by different values for n_1 and n_2 .




Figure 3.14: Duffing ($n = 1$) and higher-order ($n = \{2, \infty\}$) oscillator responses for different quality factors Q . The system modeled is an array-based Josephson parametric amplifier with $N = 20$ Josephson junctions in series. The phase response of a reflected output tone $\arg(b_{\text{out}})$ near the critical frequency and amplitude ($b_{\text{in}} = 0.98b_{\text{in,c}}$) contains all of the meaningful information, as there is no internal loss in the system. The agreement within a trio of curves (darker shade indicates higher n) of a given color (corresponding to a given Q) indicates how accurate the Duffing approximation is. The Duffing approximation is more accurate at high Q . Trios of curves are vertically offset for visual clarity. Note that with the horizontal axis expanded, all curves are seen to undergo a full 2π phase drop as expected for a tone reflected off of an internally lossless resonance.

Array-based JPAs that meet Eq.(3.198) can thus confidently take advantage of Duffing oscillator physics in their design and operation [148]. In addition, geometric inductance, L_G in Fig. 3.6a, can be used to achieve higher dynamic range without violating the Duffing approximation.⁷⁵

⁷⁵ The work of this thesis takes advantage of the equivalence between stray inductance and arrays in order to use the Duffing approximation, but not to install additional geometrical inductance. Nevertheless, the tradeoff between the compactness (and broader tunability, when SQUIDs are used) of arrays, and the more idealized behavior of single junctions with linear inductances, is of potential value for future work.

3.6 Current-pumped JPA implementation

Key results in this thesis rely on the use of current-pumped as well as flux-pumped JPAs.⁷⁶ This section discusses important design considerations and basic characterization measurements for the current-pumped JPAs used in Ch. 4, and the next section does the same for the flux-pumped JPAs used in Chs. 5 and 8. The specifications of our SQUIDs, common to both current- and flux-pumped JPAs, is included in this section. The JPAs used in this thesis were designed using Xic Graphical Input Editor from Whiteley Research, Inc and KLayout open source software.

3.6.1 SQUID array specifications

Our JPAs were fabricated at NIST Boulder using the Nb-AlO_x-Nb trilayer process [202, 203] discussed in Appx. A. The parameters that define our SQUIDs are the junction critical currents I_c , the geometric inductance L_g of the SQUID loops, and the area a of the SQUID loops for magnetic flux to penetrate.

The $2.5 \mu\text{m} \times 2.5 \mu\text{m}$ junctions were designed with a nominal critical current density of $0.8 \mu\text{A}/\mu\text{m}^2$, for a critical current of $5 \mu\text{A}$. In practice, our junctions tended to come out within 40% of this value across all fabrication runs.⁷⁷ A typical set of measurements taken on one of our flux-pumped JPA chips is shown in Fig. 3.15. The room temperature resistance measurements of Fig. 3.15a were performed on individual junctions situated within test structures. The layout for a test structure, which contains four leads in order to cancel out non-junction resistances, is shown in the inset of Fig. 3.15a. The critical currents for the same five junctions, Fig. 3.15b, are obtained from the room temperature resistances via the Ambegaokar-Baratoff relation, Eq. (3.1), with $2\Delta/e = 3.0 \text{ mV}$. For the five junctions on this chip, $I_c = 5.99 \pm 0.12 \mu\text{A}$. The $\approx 20\%$ absolute error from the nominal value corresponds to a $\approx 10\%$ change in the same direction in the resonance frequency, Eq. (3.73). This imprecision is tolerable given the frequency-tunability

⁷⁶ In addition to their uses in the work presented in this thesis, the JPAs discussed in this thesis have helped enable a range of experiments performed in the Lehnert Lab at JILA [165, 168–170].

⁷⁷ The most variation occurred between different fabrications (and hence wafers). The next most variation in I_c occurred between chips within the same wafer, and the least within a single chip.

of the devices. The much smaller 2% intra-chip variation is important for achieving good device performance [173].⁷⁸

Figure 3.16 shows the geometry of one of our SQUIDs, from which the geometrical inductance and SQUID area can be estimated. The circumference of the region ringed in by the superconductor is 20 μm . Using the rule of thumb that a 1 μm length of line corresponds to a 1.2 pH inductance, $L_g \sim 24$ pH. The nominal $I_c = 5$ μA critical currents imply $L_{J0} = 66$ pH and $L_{S0} = 33$ pH by Eqs. (3.10) and (3.33), respectively. Thus, $\beta_L \sim 0.36$ via Eq. (3.42), better than a factor of five clear of being in the flux-trapping regime, Eq. (3.43).

The area a of the SQUID loop should really be thought of as an effective area consisting of the 19 μm^2 superconductor-enclosed region, plus a fraction of order 1/2 of the cell's superconducting surface area. This is because an externally applied flux that would otherwise penetrate the superconductor will be moved off of it by the Meissner effect [172]. Knowing a very precisely is not essential, because we can within reason compensate for lower a with higher bias currents and/or mutual inductances between the SQUID loop and whatever we are using to couple flux into it. So a rough estimate of $a \sim 50$ μm^2 will suffice.

3.6.2 TKC specifications and simulations

Our current-pumped TKCs consist of a capacitor C -shunted N -SQUID array capacitively coupled (C_c) to an external environment. We do not deliberately add any stray inductance L_G , and what little is there is of secondary importance to the Josephson inductance. Figure 3.17 shows two scanning electron microscope (SEM) images of a device. The capacitors C and C_c are implemented with interdigitated-finger geometries, which avoid some of the dielectric loss of parallel plate geometries at the cost of compactness. The main capacitor C actually consists of two capacitors $C/2$ in a symmetric, parallel configuration.

To ensure broad frequency coverage, we designed eight different JPAs with bare resonance

⁷⁸ In light of the observation that I_c variation grew with the physical separation of the junctions (Footnote 77), the fact that the junctions in the test structure are much farther apart than those in a given array suggests that the variation within an array will be less than or equal to the variation measured for different test structures on a single chip.




Figure 3.15: (a) Room temperature Josephson junction resistance measurements for five different junctions on a typical chip. The inset shows the layout for the five test structure used to perform the measurements. The four leads allow us to cancel out lead resistances when measuring the junction’s resistance. Different patterns indicate different electrical layers, discussed in Appx. A (solid = base electrode, striped = counter electrode). The $2.5 \mu\text{m} \times 2.5 \mu\text{m}$ junctions are too small to see at this scale, but are situated where the two thin lines at the center of each test structure cross. (b) Critical currents for the same five junctions. The critical currents are obtained from the room temperature resistances using the Ambegaokar-Baratoff relation, Eq. (3.1).

frequencies logarithmically spaced between 4 and 12 GHz. The layouts and the target resonance frequencies are shown in Fig. 3.18. The downward-frequency tunability of each device was designed to overlap with at least one device below it. Parameter values were chosen in accordance with




Figure 3.16: Superconducting niobium wiring layout for a pair of SQUIDs. The cell shown here corresponds to Fig. A.1e. Dimensions useful for estimating the effective area of the SQUID loop and its geometrical inductance are provided. Blue (pink) corresponds to base electrode (wiring 1) layer discussed in Appx. A.




Figure 3.17: Scanning electron microscope images of a TKC with false-color highlighting the circuit elements. The capacitor C (blue) is implemented as a pair of interdigitated finger capacitors $C/2$ in parallel. The inductance L consists primarily of N SQUIDs, where N varies across different TKC designs. The waffling pattern covering the niobium ground plane pins magnetic flux trapped in the superconductor in place. A closeup view of part of the SQUID array shows four Josephson junctions, located at the ends of the bright ‘H’ shape. With the base niobium electrode (dark gray), they form two SQUID loops, as illustrated in Fig. A.1e. Figure reproduced from Ref. [204].

Eqs. (3.73), (3.76), and (3.58), targeting $Q_e = 50$ and $Z_{LC} = 50 \Omega$.

Electromagnetic simulations of the resonators shown in Fig. 3.18 are shown in Fig. 3.19.⁷⁹ Figure 3.19a shows the behavior the phase of S_{11} , the reflection coefficient off of the JPA’s only port, with frequency. The resonance frequency lies at the center of the 2π phase drop. It and the

⁷⁹ These simulations treat the SQUID arrays as linear inductors.




Figure 3.18: Layouts and design parameters for a suite of eight current-pumped TKCs. The frequencies listed are the targeted values, logarithmically spaced between 4 and 12 GHz. The number N of SQUIDs, and the design values for resonator and coupling capacitances C and C_c are listed for each device.

external coupling γ_e (Fig. 3.19c) were fitted for each device using the equation

$$\arg(S_{11}) = \text{atan2} \left(\frac{\gamma_e(\omega - \omega_{\text{TKC}})}{(\omega - \omega_{\text{TKC}}) - \gamma_e^2/4} \right), \quad (3.199)$$

where atan2 is the two-argument arctangent function. Equation (3.77) then gives the external quality factor Q_e (Fig. 3.19b). The resonator impedance Z_{LC} was determined from the individual resonator components and Eq. (3.58).

3.6.3 Characterization measurements

A basic set of measurements is performed to confirm proper JPA functionality and characterize the device properties, including the bare resonance frequency, tunability, gain, and bandwidth. The basic configuration of the TKC together with other microwave elements that realizes the current-pumped JPA is shown in Fig. 3.20. Here, a directional coupler [187] is added in order to apply a pump tone.

These measurements, as well as those in Sec. 3.7.2, require combinations of three key instruments: a microwave generator for producing pump tones, a vector network analyzer (VNA) for sending and receiving probe tones, and a DC current source for sourcing a DC magnetic flux through a superconducting coil. Most commonly, we used a E8257D PSG Analog Signal Generator




Figure 3.19: (a) Simulated responses for the current-pumped JPAs of Fig. 3.18 at their maximum resonance frequencies. The resonance frequencies $\nu_{\text{TKC}} = \omega_{\text{TKC}}/2\pi$ fitted from reflected scattering parameter phase $\arg(S_{11})$ versus probe frequency ν are indicated, along with the number N of SQUIDs for each device. (b) Simulated external quality factors Q_e for each device. (c) Simulated external coupling rates γ_e . Coupling rates increase with ν_{TKC} as a consequence of holding Q_e fixed. (d) Resonator impedances Z_{LC} obtained from simulated circuit element values. Simulations were performed using the Sonnet Software’s electromagnetic simulator.

and a E5071C ENA Vector Network Analyzer, both from Keysight Technologies, for the microwave generator and VNA, respectively, and a Yokogawa GS200 DC Voltage/Current Source to provide our DC current.




Figure 3.20: JPA configuration for current pumping. The TKC (Fig. 3.6b) is shown within the JPA (large triangle) as in Fig. 3.7. Here, details specific to biasing the JPA for current pumping are included. A directional coupler (rectangle) couples a large pump tone into the signal path. The cylinders represented segments of transmission line. Both pump and signal emerge from the output port of the current-pumped JPA.

Figure 3.21 shows measurements of the resonant frequencies for the eight different JPAs. In this measurement, only a weak signal tone is sent into the JPA to probe its linear properties (Fig. 3.21a). The results, Fig. 3.21b, are seen to match up reasonably well to the simulated values from Fig. 3.19a.

Using the same setup, we look in Fig. 3.22 at the flux tunability of one of the JPAs. Here, a current is varied through a superconducting magnetic coil placed a few cm from the SQUID loops. The mutual inductance between the SQUID loops and the coil allows an applied DC current to source an external flux Φ_{ext} which together with the loop inductance gives rise to the total flux Φ_{tot} penetrating the loop, Eq. (3.38). The total flux tunes the SQUID inductance up according to Eq. (3.28), and hence the resonance down from its peak value according to Eq. (3.73). Ideally, the peak value should occur at zero DC bias current, but it is often offset, as shown in the figure. The offset indicates the presence of some trapped flux⁸⁰ that is reasonably homogeneous over the spatial extent of the array, and is therefore not a large problem, as it can be compensated for with

⁸⁰ This is distinct from entering the flux-trapping regime of Sec. 3.2.2.




Figure 3.21: (a) State of pump and probe for measurement of resonance frequency. No pump is applied, and a probe tone is applied in the low-power limit to measure the resonance frequency of the current-pumped JPAs. (b) Reflection measurements of phase versus frequency alongside simulations for eight different current-pumped JPAs with nominal resonances spaced logarithmically between 4 and 12 GHz.

the coil. Other nonidealities include the presence of standing waves (faint horizontal streaks), and small structure⁸¹ off of the main lobes. But overall, this is a healthy JPA, and these effects did not interfere significantly with operation.

In Fig. 3.23, a current-pumped JPA is operated for two-quadrature gain. This time, we do apply the pump tone (Fig. 3.23a). Figure 3.23b shows three gain curves, with higher gains being achieved closer to the critical point, beyond which bifurcation sets in. Here we observe the preservation of the GBP, Eq. (3.133).

This concludes our introduction to the specific current-pumped TKC/JPA devices used in this thesis. We will return to them in Ch. 4.

⁸¹ An earlier fabrication produced JPAs with much more structure off of the main lobes. This turned out to be a consequence of placing the SQUID array too close to the ground plane of the TKC, causing an inhomogeneous flux environment across the SQUIDs. These JPAs were more difficult to tune, but still operable. We fixed this error on a subsequent fabrication by simply adding 50 μm of line on either side of the array. The JPA in Fig. 3.17 is a device from the problem batch: observe that its SQUID array runs right up to the ground plane on either side. This issue was distinct from the flux-tuning problem experienced by our flux-pumped JPAs, Fig. 3.32 and Appx. B.




Figure 3.22: Flux tuning of a nominal-7.5 GHz current-pumped JPA resonance with a DC current. No pump tone is applied, as in Fig. 3.21a. When overcoupled, resonance occurs at the probe frequency ν where the reflection coefficient phase $\arg(S_{11})$ drops 2π , and tunes with a DC current applied to a superconducting coil which threads a magnetic flux through the SQUID loops. The resonance is ideally periodic in current, and the two lobes closest to 0 current are displayed here. The faint horizontal streaks are standing waves in the experimental setup. The small features off to the side of the main lobes reveal slight nonidealities either in the array fabrication or the uniformity of the magnetic flux penetrating the SQUID loops.

3.7 Flux-pumped JPA implementation

Recall from Sec. 3.4.7 that flux pumping works, in an idealized picture, by creating a time-varying inductance (and thus resonance) at pump frequency ω_p in order to amplify (or deamplify) tones about $\omega_p/2$ [141, 182]. With flux pumping, this is achieved without a strong current tone near the amplification band. The flux-pump tone modulates the resonance frequency along one of the lobes in Fig. 3.22. In order to vary the resonance at ω_p , the variation, whose scale is set by the linewidth of the resonator according to Eq. (3.170), should not cross the peak, otherwise the resonance will go up then down and then up and down again over one period of pump-tone oscillation. Even adjacent to the peak, the lobe's slope is often too shallow to achieve good performance. In practice, our flux-pumped JPAs perform best when operated at least ~ 100 MHz below the top.

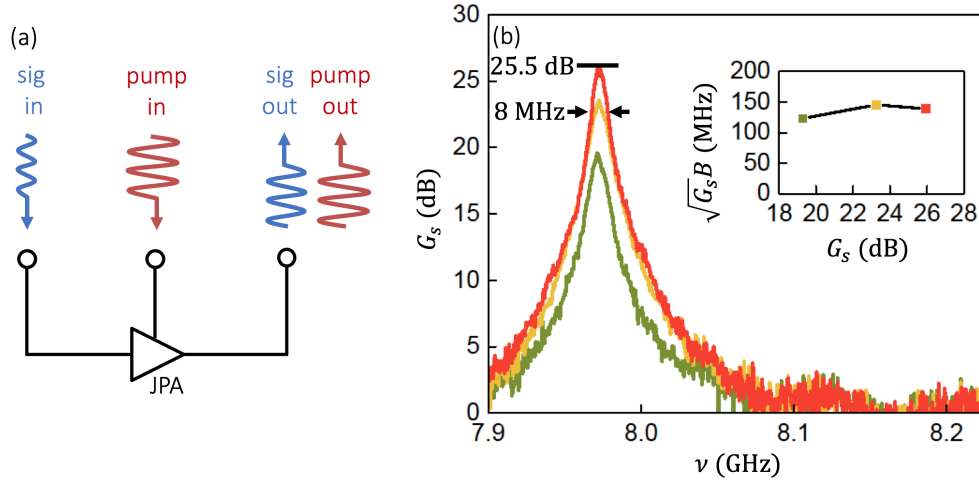


Figure 3.23: (a) State of pump and probe for measurement gain and bandwidth. A large pump tone is applied in order to amplify a small signal. (b) Gain and bandwidth measurements performed on a nominal-8.8 GHz JPA. Three different Lorentzian gain profiles show the two-quadrature JPA power gain G_s for different operating points on the same device. The bandwidth B is the full-width at half-max of the Lorentzian, and in the inset its product with the amplitude gain is seen to be a fixed quantity.

The flux- and current-pumped JPA exist to perform similar measurements, but each have distinct advantages and disadvantages. The current-pumped JPAs of the previous section are simpler, smaller⁸² devices, and easier to fabricate. They have a slightly broader tunability, as flux-pumped JPAs cannot operate all the way up to the top of the flux-tuning curve. They also do not require operation in a second frequency band a factor of two away from the band of interest. For the most part, these advantages concern the design of and hardware used to operate the JPAs. If one is willing to pay the overhead cost, flux-pumped JPAs confer significant operational advantages [141, 183–186]. First, though they require a second port on the TKC for pumping, they no longer require the lossy directional coupler of Fig. 3.20 in the signal path.⁸³ Second, the pump tone is not routed to the JPA output, and is furthermore not within the band of interest. Third, an equivalent flux-pumped JPA will have higher power handling than its current-pumped counterpart,

⁸² Smaller here refers to the TKC footprints on a given chip. Designs used in this thesis include four (two) current-pumped (flux-pumped) devices per chip.

⁸³ Current-pumped JPAs can also avoid this, but doing so requires adding on a second, weakly coupled port to the device [140] not so different than the one used for flux pumping.

and will not have the sizable Kerr-induced frequency shift of the resonance.⁸⁴ And, as a bonus, a flux-pumped JPA, if desired, can even be current pumped (see e.g. Fig. 3.35).

3.7.1 Flux-pumped JPA design

We optimized several constraints in the design of our flux-pumped JPAs.⁸⁵ First, we took care to ensure uniform coupling between the pump line and each of the SQUIDs that collectively constitute the tunable inductor. Uniform coupling guarantees that the JPA’s dynamics are spatially homogeneous, resulting in improved power handling. Second, we placed the pump line close to the SQUID array, in order to minimize the pump power required to drive the LC resonator. At the same time, the coupling of the LC circuit to the pump line is kept much lower than the coupling to the line carrying the signal near the amplification band, in order to avoid losing part of the signal through the pump port: the transmission between the two ports is kept to nearly -30 dB. Finally, we shaped the flux line as a ‘U’ around the SQUID array. With this configuration, the pump couples to the differential-mode current circulating inside each SQUID loop and is isolated from the common-mode current, unidirectional across the SQUID array.

We considered and simulated many designs, including those shown in Fig. 3.24, for flux-pumped JPAs, before settling on the relatively simple layout of Fig. 3.25a. The TKC, shown up close in Fig. 3.25b, is laid out the same as in our current-pumped JPAs of Sec. 3.6. New to the flux-pumped designs, an AC bias line runs across the length of the chip. The line tapers and goes through a small interdigitated finger capacitor $C_f \ll C_c$ acting as a choke filter to suppress low-frequency noise before wrapping around the SQUID array, crossing over the line containing the SQUIDs as it reverses direction. As shown in Fig. 3.25c, the current flowing both left and right contributes to an AC flux within the SQUID region. The bias line ends back at the ground plane to the right of the array.

⁸⁴ It is not that the Kerr nonlinearity is gone; it is just that there is no large current tone that causes the resonance to shift downward because of it. Small Kerr-induced shifts can still occur because of signals and nonideal flux-pump tones.

⁸⁵ For flux-pumped JPAs, we will frequently refer to the chip containing the TKC and flux-bias line, which is in truth neither the TKC or nor the entire JPA, as simply the “JPA.”




Figure 3.24: Discarded designs for TKCs with flux-bias lines. Six different concepts for laying out the SQUID array and flux-bias line for our flux-pumped JPAs are shown here. The shaded region indicates superconductor, and the white region dielectric.

Figure 3.26 shows simulations performed in AWR Software’s Microwave Office of important circuit behaviors which were used to guide the design of the flux-pumped JPAs. The layout and port numberings used throughout the rest of the figure are shown in Fig. 3.26a. Ports 3–5 are used to sense current along the flux bias line. Ports 6–11 are used to detect current within the SQUIDs. As discussed in Sec. 3.2, the current within the two SQUID arms can be considered in the basis of total and circulating currents. As illustrated in Fig. 3.26b, we here refer to I from Eq. (3.15) as the common-mode current, I_{com} , and recast I_{circ} as the differential current, $I_{\text{diff}} = 2I_{\text{circ}}$. It is the differential current that is essential to flux pumping, while any common current is an undesired side-effect. The AC pump current I_p flows along the bias line as in Fig. 3.25c.

The first design consideration is the uniformity of current on the bias line. Figure 3.26c shows the amplitude of the scattering parameter from port 2 to each of three internal ports spaced along the bias line.⁸⁶ The positive and negative numberings on the ports act as the positive and negative nodes of a port for the purpose of simulating the scattering matrix. The response to the

⁸⁶ For a discussion of internal ports in AWR, see Ref. [205].

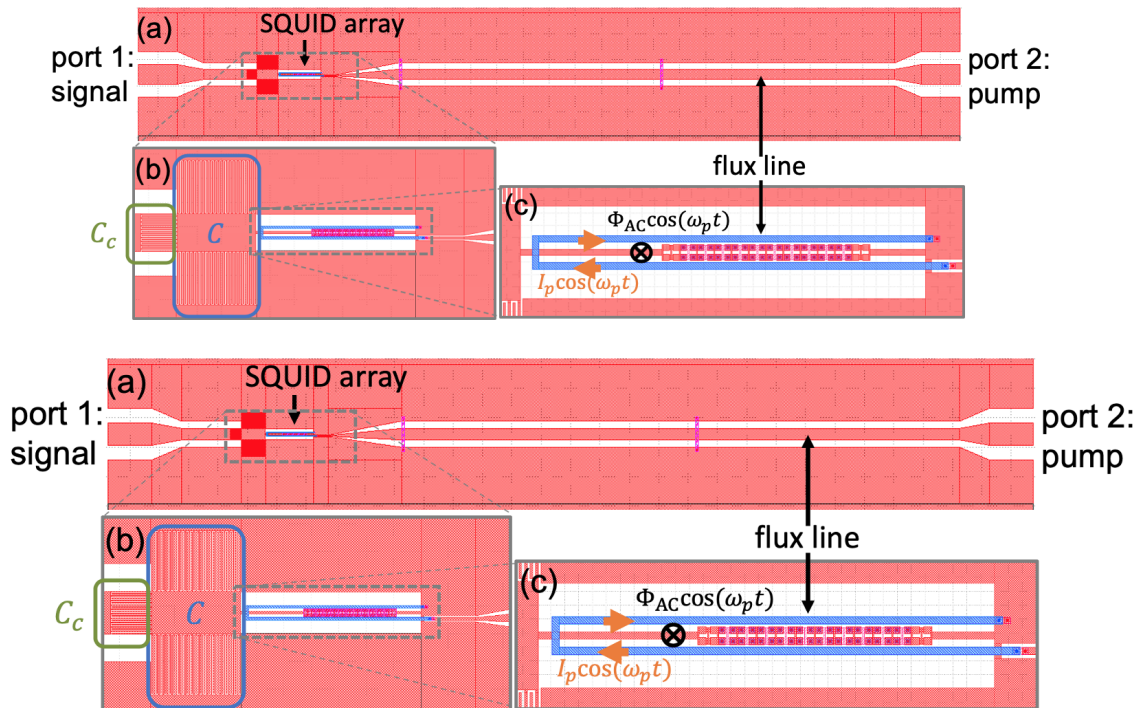


Figure 3.25: (a) Design for the full TKC connected to the flux-bias line, which extends across the entire chip (length 6.3 mm). Bridges of superconductor cross the flux-bias line to ensure that the ground plane is not interrupted. (b) Closeup on the TKC, with capacitor C (blue) and coupling capacitor C_c (green) indicated. The capacitively coupled flux-bias line enters from the right. (c) Further closeup on the flux-bias line and SQUID array. The flux-bias line meanders around the SQUID array such that the contributions from the left- and right-moving current $I_p \cos(\omega_p t)$ add according to the right-hand rule within the SQUID loops, producing an AC flux Φ_{AC} at the same frequency.

stimulus from port 2 is shown to differ from one extreme end of the array to the other by about 5 dB (difference of green and black lines; less than a factor of two in current amplitude) independent of frequency. Since the device in Fig. 3.26 is meant to operate around $\omega_{TKC} \approx 6$ GHz, we care about the behavior of currents on the flux-pump line at $2\omega_{TKC} \approx 12$ GHz. Here, we see that the even for port 3 (black curve), farthest from the flux-pump port (port 2), the transmission is greater than -10 dB. Whether this is sufficient to bias the SQUIDs for flux pumping depends on the mutual inductance between the bias line and the SQUID loops.

The second design consideration concerns this mutual inductance: in particular, we need a

sufficiently large differential current to modulate the SQUIDs with $\sim 1/Q_e$ flux quanta in order to flux pump. For $Q_e \approx 50$, Eq. (3.37) implies that we need $I_{\text{diff}} \sim 0.6 \mu\text{A}$. However, the fact that we wish to operate these devices in a cryogenic environment around $T_{\text{base}} \lesssim 30 \text{ mK}$ limits how much current I_p we can supply along the pump line. In order to reduce Johnson-Nyquist noise [206, 207] along the flux-pump port, we will couple the pump tone in with a 10 dB directional coupler (Fig. 3.29).⁸⁷ We measured the cooling power of our Oxford Triton 200 dilution refrigerator to be roughly $10 \mu\text{W}$ at 30 mK. This is the maximum power we can send to the mixing chamber plate of the dilution refrigerator, and the 10 dB coupler implies that only 10% of the power, $1 \mu\text{W}$, can make it through to port 2 of the flux-pumped JPA. For a 50Ω characteristic impedance, $1 \mu\text{W}$ corresponds to $\approx 140 \mu\text{A}$. We therefore conclude that we require enough inductance coupling to achieve

$$\left| \frac{I_{\text{diff}}}{I_p} \right| > 0.004, \quad (3.200)$$

the ratio between 0.6 and $140 \mu\text{A}$.

Figure 3.26d shows that our flux-pumped JPA meets the criterion of Eq. (3.200). By simulating the admittance matrix Y_{ij} [187] of the system, we identify

$$\frac{I_{\text{diff}}}{I_p} = \frac{Y_{63} - Y_{73}}{Y_{33}} \quad (3.201)$$

for port 3 and the adjacent SQUID. The result of Eq. (3.201) (black curve), along with analogous curves for the ports 4 (red) and 5 (green) and their neighboring SQUID loops, indicates that $I_{\text{diff}}/I_p \sim 0.025$, independent of frequency. As a confirmation on our understanding that the power that makes it to each port 3–5 is sourcing the differential current in the neighboring SQUIDs, we observe that the curves maintain the same hierarchy in Fig. 3.26d as in Fig. 3.26c.

The third and final design consideration, plotted in Fig. 3.26e, is that the common-mode current be much smaller than the differential-mode current,

$$\left| \frac{I_{\text{com}}}{I_{\text{diff}}} \right| \ll 1, \quad (3.202)$$

⁸⁷ Unlike the directional couplers used for current pumping, this one does not go in the signal path, and will therefore not be harmful to squeezing.

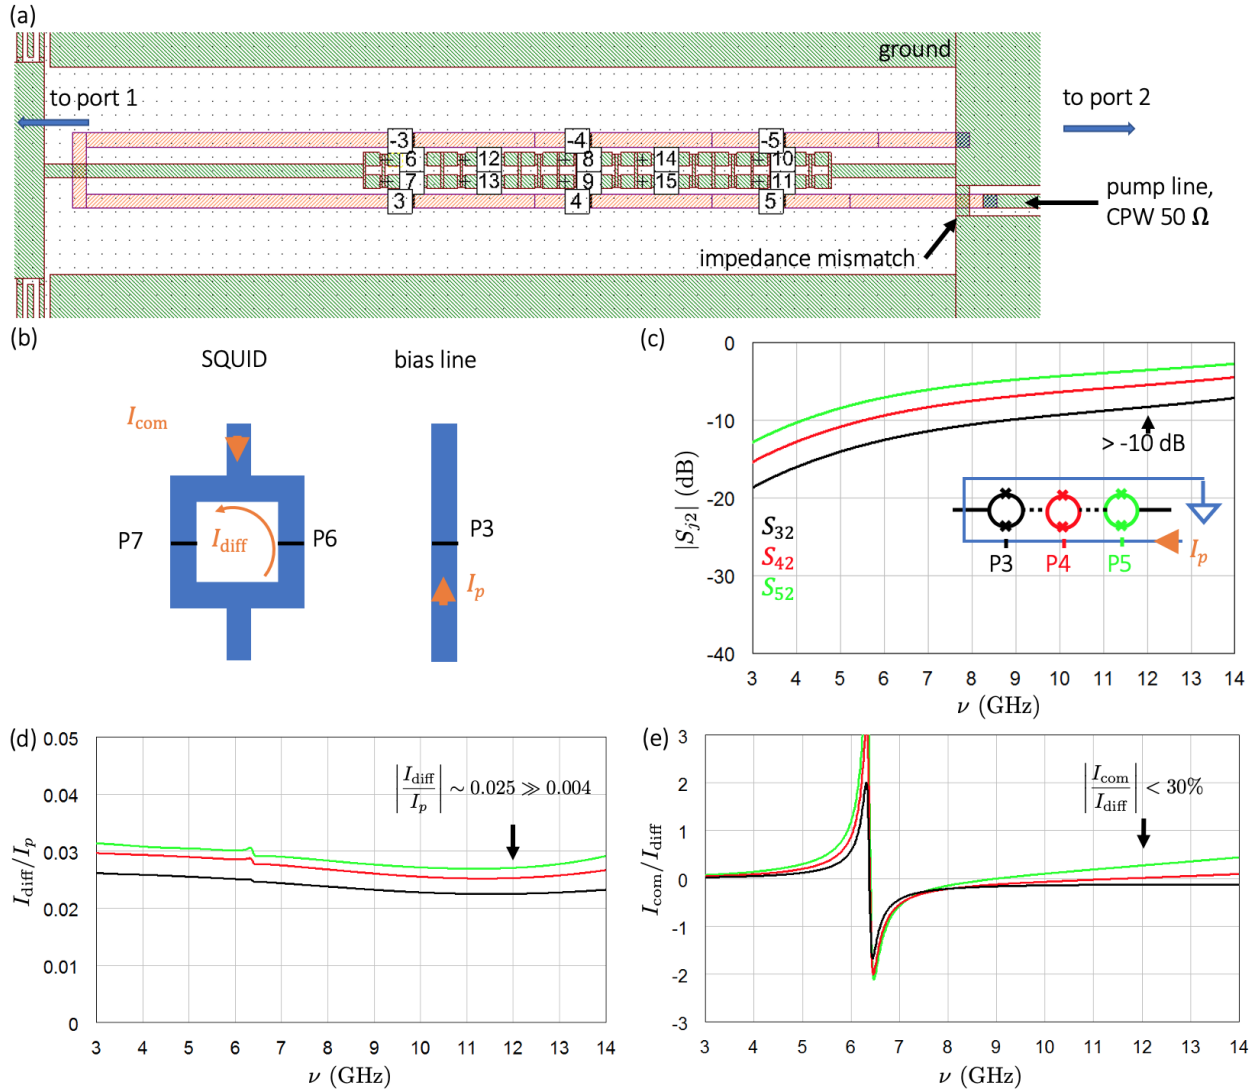


Figure 3.26: (a) Design layout for the SQUID array region of the flux-pumped JPA shown in Fig. 3.25c. Green (red) shaded region indicates the location of base electrode (wiring) superconductor. Josephson junctions are modeled with lumped element nonlinear inductances. Numbered internal ports are used to detect the current flowing at different physical locations. Ports 1 and 2 (not shown) are the input ports for signal and flux pump, respectively. Ports 3–5 measure current along the bias line. Ports 6–11 measure currents on individual arms of SQUIDs. Not all of the current sent in the 50 Ω coplanar waveguide (CPW) pump line from port 2 makes it to the bias line because of an impedance mismatch at the entry to the SQUID region. (b) Definitions of the common (I_{com})- and differential (I_{diff})-mode currents within a SQUID, as well as the pump current I_p along the AC bias line. (c) Probe frequency ν -dependent scattering parameter magnitude from port 2 to ports 3, 4, and 5, corresponding respectively to locations on the bias line near the left, center, and right of the SQUID array. (d) Ratio of differential-to-pump currents for ports 3–5 and the adjacent SQUIDs versus ν . (e) Ratio of common-to-differential currents for the SQUIDs adjacent to ports 3–5 versus ν . For (c)–(e), we are chiefly interested in the behavior around 12 GHz for this flux-pumped JPA, designed to work near 6 GHz.

when port 2 is excited. Failing to meet Eq. (3.202) would mean a substantial fraction of our pump tone would be wasted creating a common-mode current, useless for flux pumping and able to partially escape out the signal port of the TKC. For port 3,

$$\frac{I_{\text{com}}}{I_{\text{diff}}} = \frac{Y_{63} + Y_{73}}{Y_{63} - Y_{73}}. \quad (3.203)$$

Equation (3.203) (black curve) and its analogues for ports 4 (red) and 5 (green) with neighboring SQUIDS show that the differential current for $\omega_p = 12$ GHz is at least threefold larger than the common current across the array.⁸⁸




Figure 3.27: (a) Simulated reflection coefficient phase $\arg(S_{11})$ versus probe frequency ν for the signal port (port 1) of the flux-pumped JPA of Fig. 3.26. (b) Simulated transmission coefficient magnitude $|S_{21}| = |S_{12}|$ versus ν between the signal and pump port (port 2) of the same JPA.

Satisfied that the flux-bias line behaves as desired, we now turn our attention to the behavior of the entire, two-port TKC. Figure 3.27a shows the reflection coefficient phase $\arg(S_{11})$ off of the signal port, indicating $Q_e \approx 50$, as desired. Figure 3.27b shows transmission between the signal and pump ports. We are concerned both with the behavior at the signal ~ 6 GHz and pump ~ 12 GHz frequency. Near the pump frequency, transmission is below -40 dB, which is practically ideal behavior. On resonance, transmission reaches -28 dB, indicating that only 0.15% of a tone sent into the signal port will be lost out of the flux-pump port.

Finally, using the harmonic balance functionality in Microwave Office, we verify in Fig. 3.28 that the device can achieve gain if the SQUIDS are properly biased with both DC and AC currents.

⁸⁸ Since Microwave Office does not account for the Meissner effect, we might expect the precise values in these plots to change slightly when the environment seen by the magnetic flux is partially superconducting.




Figure 3.28: Simulated gain G_s versus probe frequency ν curves for different AC bias strengths using harmonic balance. Pump powers used range from -55 to -48 dBm input to the flux-pump port. Probe power is set to -120 dBm.




Figure 3.29: JPA configuration for flux pumping. This figure is the specification of Fig. 3.7 and analogue of Fig. 3.20 for flux pumping. The directional coupler is now outside of the signal path (and is not considered part of the JPA), and the pump tone no longer emerges with the signal.

We do not apply the DC currents through the port 2 bias line (which would not work in any case due to its capacitive coupling), but rather manually by equipping each of the SQUIDS with a DC current bias that creates a circulating current.

We designed six flux-pumped JPA chips with bare resonances logarithmically spaced between

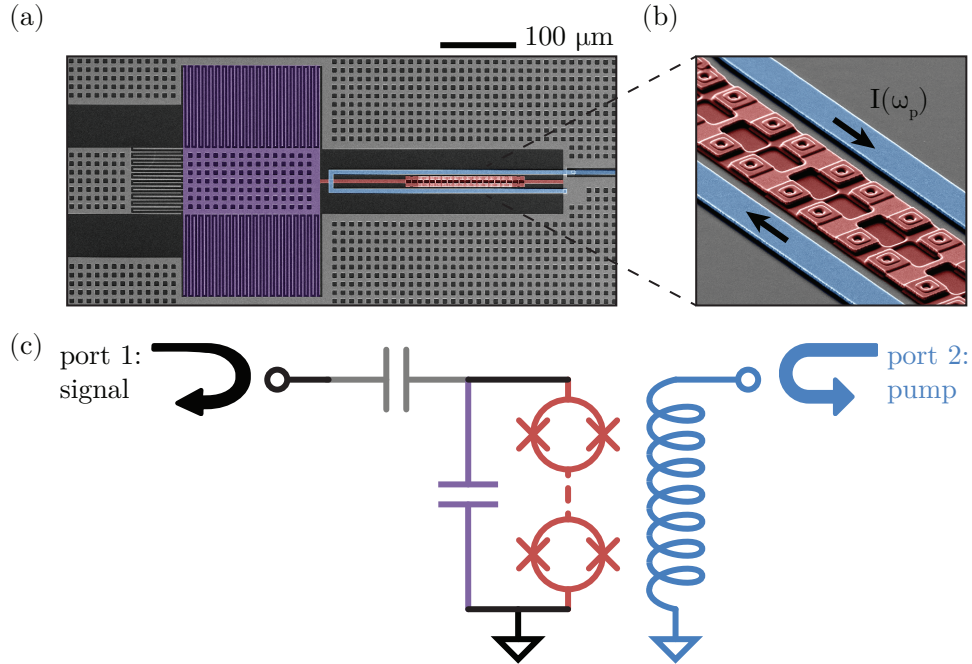


Figure 3.30: (a) False-color scanning electron microscope image showing the main elements in one of our flux-pumped JPAs. Elements include a 550 fF interdigitated finger capacitor (purple), an inductance comprising a SQUID array of 6 μA Josephson junctions (red), and a flux line, shorted to ground (blue). (b) A closeup view shows several SQUID loops (red), surrounded by the flux line (blue). The pump current $I(\omega_p)$ sources a differential-mode current around the SQUID loops. (c) The JPA equivalent circuit is shown with the same color-coding used in the images of (a) and (b). Figure reproduced from Ref. [208].

4.6 and 10.3 GHz.⁸⁹ Next, we will characterize their performance with some basic tests, as we did for the current-pumped JPAs in Sec. 3.6.3.

3.7.2 Characterization measurements

The configuration of the TKC with other microwave components to realize a flux-pumped JPA is shown in Fig. 3.29. Two of the advantages of flux-pumped JPAs — the removal of the directional coupler from the signal path and the absence of the flux-pump tone at the output — are visible here. False-color, SEM images are shown in Fig. 3.30 together with the equivalent circuit for comparison. The JPA is pictured mounted and wire-bonded to a printed circuit board in Fig. 3.31.

⁸⁹ This is essentially the same set of frequencies as for the current-pumped JPAs of Sec. 3.6.2, but leaving off the lowest (4 GHz) and highest (12 GHz) frequency chips.




Figure 3.31: Photo of a mounted flux-pumped JPA. The JPA is wire-bonded to a printed circuit board and attached with Elmer’s rubber cement to a gold-plated copper mount. The assembly is shown next to a penny for scale. It is later in a light-tight box, the bottom of which is formed by the mount, and interfaced with via SMP connectors.

Figure 3.32 shows a two-dimensional (bias current, frequency) sweep of the bare JPA phase-response with a weak probe tone and no pump. This measurement is identical to that of Fig. 3.22 for a current-pumped JPA. The immediate difference between the two is obvious: the flux-pump tuning curve is substantially more structured. This structure is undesirable, and should not be present in flux- or current-pumped designs. Its source remained a mystery over the course of the experiments carried out with flux-pumped JPAs in Chs. 5 and 8, but was ultimately discovered to be trapped flux in the “dummy SQUID” loops placed at the ends of the SQUID array. The problem, and its resolution, are discussed in Appx. B. While the additional structure made tuning somewhat more difficult, it generally did not limit performance once a satisfactory tuning location was found.

In Fig. 3.33, we apply a flux-pump tone and successfully realize gain. We can modestly (~ 50 MHz) tune the device’s amplification band at a given DC flux bias by changing the half-pump frequency $\omega_p/2$. To achieve more dramatic tuning, the coil bias current is tuned.

The broader tuning landscape in the space of half-pump frequency and pump power is shown in Fig. 3.34a for a given flux-bias. For linear amplification, the JPA should be operated *below* the




Figure 3.32: Flux tuning of a nominal-7.5 GHz flux-pumped JPA resonance. This figure shows the same measurement of reflection phase $\arg(S_{11})$ versus probe frequency ν and DC bias current that Fig. 3.22 shows for a current-pumped JPA. The tuning landscape is here significantly more messy, because dummy SQUIDs placed on the end of the SQUID arrays to create a more homogeneous superconducting environment around each SQUID had the unintended effect of trapping flux near the arrays.

point where gain diverges (or would diverge, in the small-signal limit) at a given frequency — that is, below the center of the bright red feature in Fig. 3.34a. Figure 3.34b–d show gain curves taken at various powers along the three dashed lines in Fig. 3.34a. Better operation tends to occur closer to the line of 3.34b, where the high-gain region bottoms out in pump power.

Finally, Fig. 3.35 shows the 1-dB compression point for the flux-pumped JPA as a characterization of the power handling, and compares it to that for current pumping on the same device. The 1-dB compression is defined as the signal power where the gain falls off by 1-dB from its low-power value. For higher gains, we observe that it takes smaller signals to saturate the power handling of the device. However, we see that flux pumping outperforms current pumping on the same device




Figure 3.33: Flux-pumped JPA gain G_s versus probe frequency ν curves showing tunability between and within different coil tunings. Within two DC flux tunings ~ 300 MHz apart in resonance frequency, G_s curves with maximum gain over 30 dB are achievable across ~ 50 MHz.

for all gains.

With a knowledge of both current- and flux-pumped JPAs in hand, we are now ready to move onto several key experimental results of this thesis. We will begin in the next chapter with an experimental technique to improve the power handling for current-pumped JPAs used for squeezing.




Figure 3.34: (a) Color map showing flux-pumped JPA two-quadrature gain G_s in the two-dimensional parameter space of pump power and half-pump frequency $\nu_p/2 = 2\pi\omega_p/2$. Pump powers are estimated at the input of the JPA flux-pump port from. (b)–(d) Gain versus probe frequency ν curves taken along the linecuts indicated in (a). Some curves taken higher in power than the main feature in (a) at a given frequency sometimes show a self-oscillation spike at the half-pump frequency.




Figure 3.35: Power-handling measurements for flux- and current-pumped JPA operation. The 1-dB compression point input signal power decreases with gain G_s for both flux- and current-pumped operation of the same nominal 7.5 GHz flux-pumped JPA. Flux-pumped operation typically achieves ~ 5 dB higher power handling for this device.

Chapter 4

Optimal current-pumped Josephson parametric amplifier operation for vacuum squeezing

When operated in a phase-sensitive mode, the Josephson parametric amplifiers (JPAs) discussed in the previous chapter were seen to amplify one quadrature of the electromagnetic field while deamplifying the other, enabling the squeezing of vacuum states below the standard quantum limit. Such squeezed states can improve qubit readout fidelity in dispersive measurements [209–212], help to better measure the motion of a mechanical oscillator [213], and enhance the signal-to-noise ratio of spin-echo detection in a magnetic resonance experiment [214]. Relevant to this thesis, squeezing can also accelerate haloscope-based searches for axionic dark matter. Where we left off in Ch. 2, JPAs had already been successfully deployed in haloscopes, albeit only in a phase-insensitive mode of operation [90, 91, 106, 107].

In the idealized picture of Sec. 3.4.6, a JPA squeezes one quadrature of vacuum while amplifying the other by the same amount [140, 215, 216]. Unfortunately, in practice vacuum squeezing is limited by microwave losses and distortion due to amplifier saturation, or finite pump stiffness [179, 217]. Minimizing loss through improved, low-loss elements or via fewer microwave connections remains an important challenge, and, to that end, on-chip circulators [218–222], flux-pumped JPAs (Sec. 3.7), and directional amplifiers [149, 151, 218] are promising. In this chapter, we will focus on the complimentary challenge challenge: distortion.

In Ch. 3, we saw that a JPA is essentially a nonlinear LC resonator whose nonlinearity comes from the Josephson junctions comprising its inductance. Parametric oscillation of the circuit

converts a pair of pump photons at frequency ω_p into a signal photon at ω_s and a complementary idler photon at ω_i in a four-wave process: $2\omega_p = \omega_s + \omega_i$. In the degenerate case, the pump is centered about the amplification band and $\omega_p \approx \omega_s \approx \omega_i$.

Different combinations of pump power and frequency will produce the same gain, and one commonly operates at points first order-insensitive to pump power fluctuations, i.e. on the line of maximum gain (LMG) [181]. Here a typical JPA can achieve more than 20 dB of gain and maintain a linear relationship between input and output signal powers. Squeezing, however, is more susceptible to nonlinearities that distort the output, arising when signal photons in the resonator greatly affect the pump's amplitude. These nonlinear effects impact the squeezed quadrature at gains much lower than those where they affect the amplified quadrature. Pump stiffness can nonetheless be improved by double pumping or flux pumping at the cost of increased design and control complexity.

In this chapter, which presents the work of Ref. [204] adapted to the language, conventions, and context of this thesis, we will see how a current-pumped JPA can be optimized for squeezing. In particular, pumping it with a power greater than that which maximizes gain at a given pump frequency improves pump stiffness and minimizes signal distortion. We will first experimentally study this effect on a coherent tone whose power is comparable to vacuum, and measure how it is transformed by a JPA operated at constant gain, for various pump powers and frequencies. Then, we will determine the optimal operating point for which the distortion is minimized. Finally, using a second JPA as a phase-sensitive preamplifier to measure a vacuum squeezed state prepared by the first JPA, we will see that squeezing follows the same dependence on operating point. With optimized pump power and frequency, and while keeping microwave losses low, this technique will allow for the generation, transportation, and readout of a state 3.87 ± 0.03 dB below vacuum, over a bandwidth of 30 MHz, in an experimental setup designed to be equipped to an axion haloscope.

We will see that we are able to improve further on these numbers in the next chapter by switching to flux-pumped JPAs. Nonetheless, the technique for optimizing current-pumped JPAs for squeezing presented in this chapter may be of use to some axion haloscopes, as current-pumped JPAs are often easier to design and/or acquire.

4.1 Signal transformation in phase space

In Sec. 3.4.4, we derived the input-output relation between small input signal and output signals \hat{a}_{in} and \hat{a}_{out} reflecting off of a tunable Kerr circuit (TKC; Fig. 3.6). Here, we will consider just the classical behavior of the system.

A classical limit of the input-output theory derivation in Sec. 3.4.4 gives the complex field amplitude, or phasor, α_{out} , at the TKC output from the incoming field amplitude α_{in} as

$$\alpha_{\text{out}} = \alpha_{\text{in}} \left(1 - \frac{\gamma_e}{i(Kn_\alpha - \Delta) + \gamma_e/2} \right). \quad (4.1)$$

Equation (4.1) follows as an immediate consequence of Eq. (3.106) given the boundary condition Eq. (3.100). The average photon number n_α , detuning Δ , Kerr constant K , and power decay rate γ_e are given in earlier sections by Eqs. (3.110), (3.108), (3.96), and (3.77), respectively. This semiclassical model does not properly treat the quantum noise at the JPA's input, but it nonetheless makes useful predictions of its behavior on a qualitative level.

In Sec. 3.4.4 we treated the JPA as losslessly amplifying an infinitesimally weak signal $a_{\text{in}}e^{i\theta}$ in the presence of a large pump tone α_{in} (here taken as phase reference), such that

$$A_{\text{in}} = \alpha_{\text{in}} + a_{\text{in}}e^{i\theta}, \quad (4.2)$$

where we are treating both pump and signal classically. Equation (4.1) provides a mapping between input and output within the bandwidth of the JPA, and remains valid even without the stiff pump approximation (Footnote 55). This transformation being unitary below bifurcation, the gain derives from the sharp dependence of the output phase on the input amplitude. The photon number n_α is determined by the cubic Eq. (3.107). Figure 4.1a shows how a family of real input pump tones with varying amplitude α_{in} maps onto complex outputs α_{out} . They are calculated from Eqs. (4.1) and (3.107), with Δ constant. A particular input phasor and its corresponding output are indicated with black arrows, illustrating the JPA's behavior at the greatest phase sensitivity, i.e. at maximum gain.

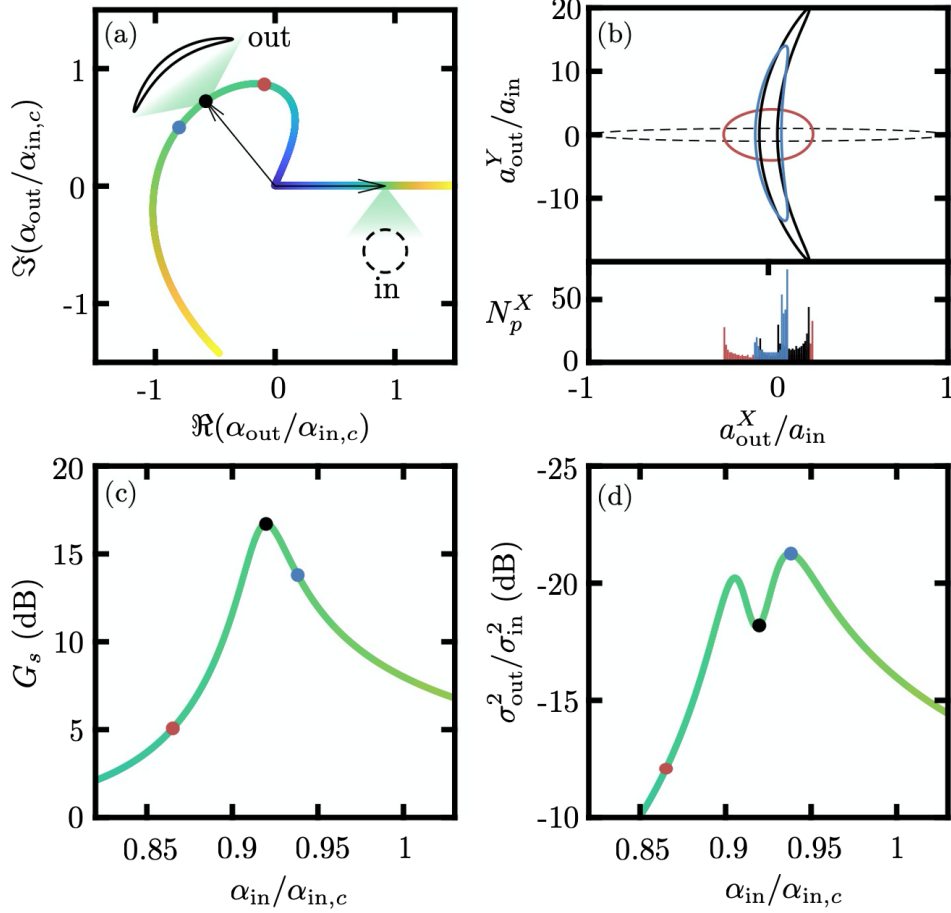


Figure 4.1: (a) Input-output mapping. A family of input pump phasors of different amplitudes describes the horizontal, colored line. These inputs map onto output phasors distributed along the correspondingly colored curve in the complex plane. They are normalized by the critical amplitude $\alpha_{\text{in},c} = \gamma_e/\sqrt{3\sqrt{3}|K|}$. Two black arrows indicate the pair of input and output phasors at maximum gain. At this operating point, a zoom-in shows how a circular input signal phasor distribution maps onto a banana at the output. Two other operating points (red and blue) are indicated. (b) Signal distortion in phase space. The input phasor, simulated with 360 different angles θ (dashed circle, deformed due to the unequal axis scaling), is processed by the Josephson parametric amplifier (JPA) for three different pump powers corresponding to the operating points from (a). The lower panel shows the distribution of points N_p^X of these phasors, when projected onto the deamplified X -axis. (c) JPA direct gain and (d) deamplified quadrature variance plotted as a function of pump amplitude. All calculations are for $\gamma_e = 2\pi \times 109$ MHz, $\Delta/\gamma_e = -0.77$, $K/\gamma_e = -4.2 \times 10^{-4}$, and a quality factor $Q = \omega_0/\gamma_e = 65$, typical values for our array-based JPAs. Figure adapted from Ref. [204].

When adding a signal $a_{\text{in}}e^{i\theta}$, the total input phasor A_{in} undergoes a θ -dependent transformation. One quadrature of the output signal a_{out} is amplified while the orthogonal one is deamplified.

At low JPA gain or infinitesimal input signal, a_{out} exhibits an elliptic shape. However, as the input signal or gain increases, it distorts into a “banana,” indicating that the pump is no longer stiff. Figure 4.1b represents such distortions, or deviations from the elliptical output of an ideal degenerate parametric amplifier [179, 214], at three different JPA gains obtained by varying the pump’s amplitude, for a probe tone whose power is equivalent to half a photon over the JPA’s bandwidth. They are rotated to share their amplified and deamplified quadratures, respectively a_{out}^Y and a_{out}^X . At the point of maximum gain, a_{out} actually has a larger projection on the deamplified axis X than at intermediate gain, where the projection onto this axis is minimal. When a signal is detuned from the pump, it creates output rotating phasors such as these, and the direct, or signal (see Footnote 61), gain is expressed as [176]

$$G_s = \frac{\langle G_\theta \rangle + 1}{2}, \quad (4.3)$$

with $\langle G_\theta \rangle$ the average ratio of output to input rotating phasor power. Figure 4.1c shows the direct gain as a function of pump amplitude α_{in} .

A vacuum state, considered semi-classically as a statistical ensemble of phasors, should undergo similar distortion when squeezed by a JPA. This effect degrades the maximal achievable squeezing, because it increases the noise variance along the deamplified quadrature axis. To make a qualitative estimate of the attainable degree of squeezing, we first project the distributions of input and output points onto the deamplified quadrature axis (Fig. 4.1b), and then compare the variances, (σ_{in}^2) and (σ_{out}^2) , of these projections. Figure 4.1d shows the ratio $\sigma_{\text{out}}^2/\sigma_{\text{in}}^2$, as a function of pump power, for a given pump frequency. Deamplification is maximized when this ratio is minimized. Strikingly, the pump amplitude that produces maximum gain fails to yield maximum deamplification due to phasor-distortion of the vacuum-sized input. Instead, deamplification reaches a global maximum at an amplitude greater than that at maximum gain, i.e. above the LMG, where phasor distortion is better mitigated. We will examine this behavior experimentally in Sec. 4.2, enabling the optimization of vacuum squeezing discussed in Sec. 4.3.

4.2 Coherent state deamplification

The first experiment discussed in this chapter investigates the distortion of a vacuum-sized coherent tone when varying the JPA operating point by changing the pump power and frequency. A single cryogenic experimental setup is used for studying distortion and vacuum squeezing, discussed in Sec. 4.3. It comprises two identical JPAs (Sec. 4.4.1) connected to each other, as shown in Fig. 4.2. The first we label the squeezer (SQ), and the second the amplifier (AMP). Each can be independently tuned via its own superconducting coil, and a variable temperature stage (VTS) is used to inject known thermal noise into the amplifier chain and measure the SQ added noise (Sec. 4.4.3). Microwave generators are used both to pump the two JPAs and to send coherent tones to the input of the chain. An in-phase/quadrature (IQ) mixer demodulates the output quadratures with respect to a local oscillator (LO). In this first experiment, the AMP is not used; tuned more than 1 GHz below the SQ band, it acts as a passive element of unit gain.

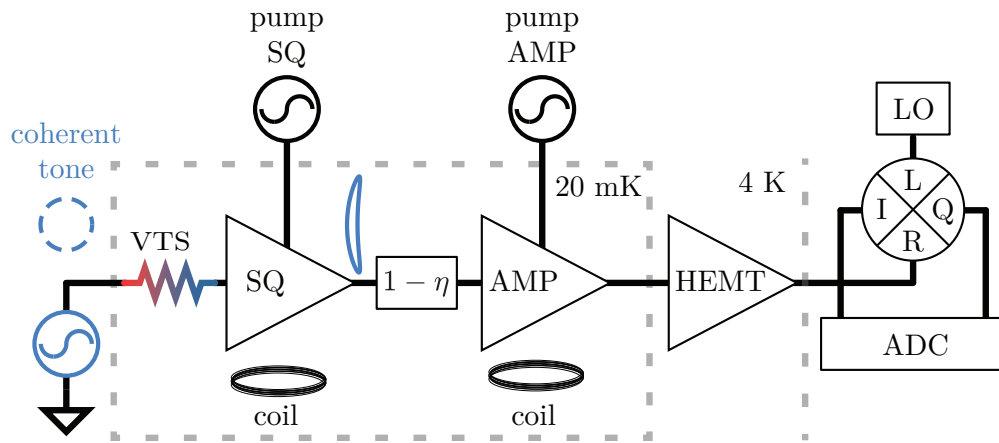


Figure 4.2: Experimental setup for coherent tone distortion and vacuum squeezing. The squeezer (SQ) and (AMP) are mounted in series. When studying distortion, the AMP is detuned and not pumped, while a microwave generator sends a known coherent tone at the chain’s input (dashed circle), then processed by the SQ. Microwave losses between SQ and AMP are parametrized by η . After subsequent amplification by a cryogenic high-electron-mobility transistor (HEMT) amplifier the tone is mapped to low frequency using an in-phase/quadrature (IQ) mixer and a local oscillator (LO). An analog-to-digital converter (ADC) records the output signal’s I and Q quadratures. Figure reproduced from Ref. [204].

In order to compare distortion and SQ direct gain G_S at particular operating points, we find

G_S (measured with a vector network analyzer, see Sec. 4.4.1) as a function of pump power P_p and frequency ν_p . We identify the critical point, Eqs. (3.117)–(3.119), at which the gain diverges, and describe the pump power and frequency as fractions of these critical values (Fig. 4.3a). The dashed black line shows the LMG. We compare deamplification when operating the SQ at points with equal gains below, above, and on this line. From the weaker dependence of gain on incident power above relative to below the LMG, visible in Fig. 4.3a, we anticipate more linear behavior above the LMG.

To characterize the nonlinear behavior and visualize phasor distortion, we now input a vacuum-sized coherent tone $a_{\text{in}}e^{i\theta}$ whose phase θ slowly rotates in the SQ pump’s frame (Sec. 4.4.1). With LO and SQ pumps at the same frequency, the tone describes traces at the output in the IQ-plane such as those in Fig. 4.3b. They are obtained with the SQ operated above the LMG. Thus, decreasing the pump power, G_S increases from 6 to 13 dB, as the output transitions from ellipse to banana.




Figure 4.3: (a) SQ gain map, with critical frequency $\nu_c = 7.0032$ GHz and input critical power $P_c = -93.1$ dBm. The different symbols \blacktriangle , \blacksquare , and \bullet correspond to those in Fig. 4.4. (b) Output tone in the IQ-plane (input is a circle of unit radius, not shown). The SQ direct gain is varied between 6 (blue) and 13 dB (green), and operated above the line of maximum gain (LMG) at $\nu_p/\nu_c = 1.0015$, a frequency indicated by arrows in (a). The asymmetric axis scales mean all output tones are longer along the Y dimension. Figure adapted from Ref. [204].

Upon repeating the measurement with the SQ off, we characterize deamplification as a function of the SQ operating point by forming the ratio $\sigma_{\text{out}}^2/\sigma_{\text{in}}^2$, as seen in Sec. 4.1. Figure 4.4a shows these ratios, calculated from the points in Fig. 4.3b. When operating above (relative to on or below) the LMG, we observe the highest deamplification at fixed gain. Above $G_S = 9$ dB, however, deamplification decreases due to phasor distortion.

Finally, we study deamplification along a contour of constant direct gain, encircling the critical point, to look for an optimal operating point at which distortion is minimized. Figure 4.4b characterizes the tone's deamplification along the $G_S = 8$ dB contour, at operating points away from bifurcation (i.e. for $\nu_p > \nu_c$), indicated in black in Fig. 4.3a. Below the LMG $\sigma_{\text{out}}^2/\sigma_{\text{in}}^2$ is about -7 dB, whereas above it decreases to about -10 dB, as a result of reduced distortion. Thus, operating the SQ above the LMG and with a gain of about 9 dB maximizes deamplification. We will leverage these findings in Sec. 4.3 to optimally bias the SQ.




Figure 4.4: (a) Deamplification shown as a function of SQ direct gain. The SQ is operated below (\blacktriangle) and above (\blacksquare) the LMG (at $\nu_p/\nu_c = 1.0015$), as well as on (\bullet) the LMG (where the pump frequency has been varied). (b) Deamplification shown for the black points in Fig. 4.3a, where $G_S = 8$ dB. Figure adapted from Ref. [204].

4.3 Optimizing vacuum squeezing with a current-pumped JPA

Most applications of microwave squeezed light require the ability to both generate and efficiently detect a squeezed state. These dual capabilities can be demonstrated with the AMP cascaded after the SQ. Upon transport from SQ to AMP, the state suffers from microwave loss η (Fig. 4.2), which replaces part of the state with unsqueezed vacuum [223], diminishing the measurable degree of squeezing. These losses were reduced to $\eta = 1.2 \pm 0.2$ dB (Sec. 4.4.2) using superconducting cables and a narrowband triple-junction circulator with insertion loss < 0.5 dB.

To make a homodyne measurement of vacuum squeezing (see full setup in Fig. 4.7, Sec. 4.4.1), one power generator provides the AMP pump and LO tones, while another, 20 kHz detuned, provides the SQ pump tone. These spectrally adjacent pump tones must be cancelled to avoid AMP saturation and preserve the squeezing of the pure vacuum state generated by the VTS at $T_{\text{VTS}} = 55$ mK $\ll \hbar\omega_S/k_B$. These cancellations are performed by combining π -phase-shifted pump tones at the output of each JPA. The AMP acts as a nearly noiseless (Sec. 4.4.3) phase-sensitive amplifier with roughly 30 dB of single-quadrature gain so as to overwhelm the rest of the amplification chain's added noise. After amplification, the squeezed vacuum is then digitized. A full 2π phase difference θ between SQ and AMP pumps is swept every $1/20$ kHz = 50 μ s, and the phase for which the digitized voltage signal is periodically minimized is taken to be the SQ squeezed quadrature. The variance of the voltage fluctuations in this quadrature is thus computed with the SQ on, and then again with it off. The reduction in variance when the SQ is on is taken to be the squeezing,

$$S = \frac{\sigma_{\text{on}}^2}{\sigma_{\text{off}}^2}. \quad (4.4)$$

Optimal squeezing occurs at $\nu_p/\nu_c = 1.0015$, with direct gain $G_S = 8$ dB, and above the LMG, as expected from Sec. 4.2. Figure 4.5a shows a histogram of the voltage fluctuations V_θ retrieved along the AMP amplified quadrature as a function of θ . Compared to panel (b), where the SQ is off, vacuum has been squeezed for $\theta = \pi/2$ and $3\pi/2$. Figure 4.5c presents a cut in the histograms at $\theta = \pi/2$, demonstrating a clear reduction in the Gaussian's standard deviation, and Fig. 4.5d displays the squeezing S as a function of θ , with a minimum of -3.87 ± 0.03 dB. This

reduction in variance, which is measured over a full 30 MHz band centered on the AMP's pump, is not a value inferred at the output of the squeezer, but rather one directly measured⁹⁰ at the terminus of the measurement chain without any losses or added noises calibrated away.




Figure 4.5: (a) Histogram of output voltage fluctuations V_θ along the AMP amplified quadrature with SQ on as a function of the phase difference θ between SQ and AMP pumps. (b) Histogram with SQ off. (c) Profile of the fluctuations and Gaussian fits (solid lines) along the dashed lines drawn in panels (a) and (b), at $\theta = \pi/2$. (d) Squeezing, Eq. (4.4), as a function of θ . The SQ is operated with direct gain $G_S = 8$ dB, obtained with a pump frequency $\nu_p/\nu_c = 1.0015$ and pump power $P_p/P_c = 0.5045$ dB. Figure adapted from Ref. [204].

Moving the SQ operating point, it was verified that squeezing decreases as we approach the

⁹⁰ This directly measured value should *not* be compared straightforwardly to inferred values of squeezing such as those cited in e.g. Refs. [224, 225]. The directly measured reduction in variance suffers from loss and added noise that comes after the creation of the squeezed state. To see the consequences of our choice, consider Ref. [215]. There, the directly measured squeezing of -1.7 dB corresponds to -9.2 dB inferred squeezing upon correcting for microwave losses and detection efficiencies. The directly measured value most closely tracks the realizable benefit of using squeezed states in the presence of loss and added noise.




Figure 4.6: (a) Squeezing plotted as a function of pump power for $\nu_p/\nu_c = 1.0015$. The points in red, black, and blue are, respectively, below, on, and above the LMG. (b) Squeezing plotted as a function of SQ direct gain. The red and blue points are those used in (a). For the black points, ν_p has been varied to obtain gains between 6 and 14 dB. Figure adapted from Ref. [204].

LMG, due to increasing gain. Adjusting the pump’s power while keeping the frequency constant, we can move the operating point from below to above the LMG, describing a line cut in the gain map, indicated by arrows in Fig. 4.3a. Figure 4.6a shows the squeezing along this cut, in qualitative agreement with the theoretical and experimental results on coherent tones seen in Secs. 4.1 and 4.2. Furthermore, a slightly better squeezing is obtained above the LMG than below, as expected. This feature is displayed in Fig. 4.6b, where squeezing is reported as a function of SQ gain.

4.4 Experimental details and calibration

This section describes additional experimental details and calibration measurements relevant to the experiments of Secs. 4.2 and 4.3. Section 4.4.1 overviews the experimental setup, and Secs. 4.4.2 and 4.4.3 respectively detail measurements to determine the loss in the lines between the JPAs and the noise added by the amplification chain.

4.4.1 Experimental setup

The full experimental setup used throughout this chapter is represented in Fig. 4.7. SQ and AMP TKCs like that of Fig. 3.17 are individually magnetically shielded. They can be tuned via coils biased by current sources. The generator which provides the AMP pump and cancellation tones also sources the LO. A switch may be used to turn off the AMP pump without interrupting the LO. A third generator (far left in Fig. 4.7) is used to create the probe tone. Output signals are demodulated with an IQ mixer, then digitized. A vector network analyzer (VNA) is used to determine the direct gains and a spectrum analyzer (SA) monitors pump cancellation. Each cancel line consists of a variable attenuator and a phase shifter. Several isolators and 3 dB attenuators are placed in the setup to minimize harmful reflections.

At base temperature, a triple-junction circulator links the SQ and AMP TKCs. It provides > 50 dB isolation between SQ and AMP, and has low insertion loss (< 0.5 dB). A VTS, weakly coupled to the base temperature plate and composed of a 50Ω load, a heater, and a thermometer, generates the chain's input noise and allows for measurement of SQ added noise. Superconducting Nb/Ti lines are placed between TKCs and directional couplers to further reduce microwave loss, and the directional couplers are directly connected to the circulator with swept elbows.

To measure the distortion of a coherent tone (Sec. 4.2), its power is adjusted to $P_{\text{PR}} = \frac{1}{2} \hbar \omega B$ in order to keep it comparable to vacuum, with B the SQ bandwidth. With the LO generator at the same frequency as the SQ pump, the tone is detuned by $\nu_{\text{PR}} - \nu_p = 20$ kHz, so that its phasor slowly rotates in the LO's frame, enabling the measurement of 360 independent phasors within the 15 MHz intermediate frequency (IF) bandwidth of the readout chain. Due to the feeble input signal, the two output quadratures must be heavily averaged in order to describe traces in the IQ-plane such as those in Fig. 4.3b.

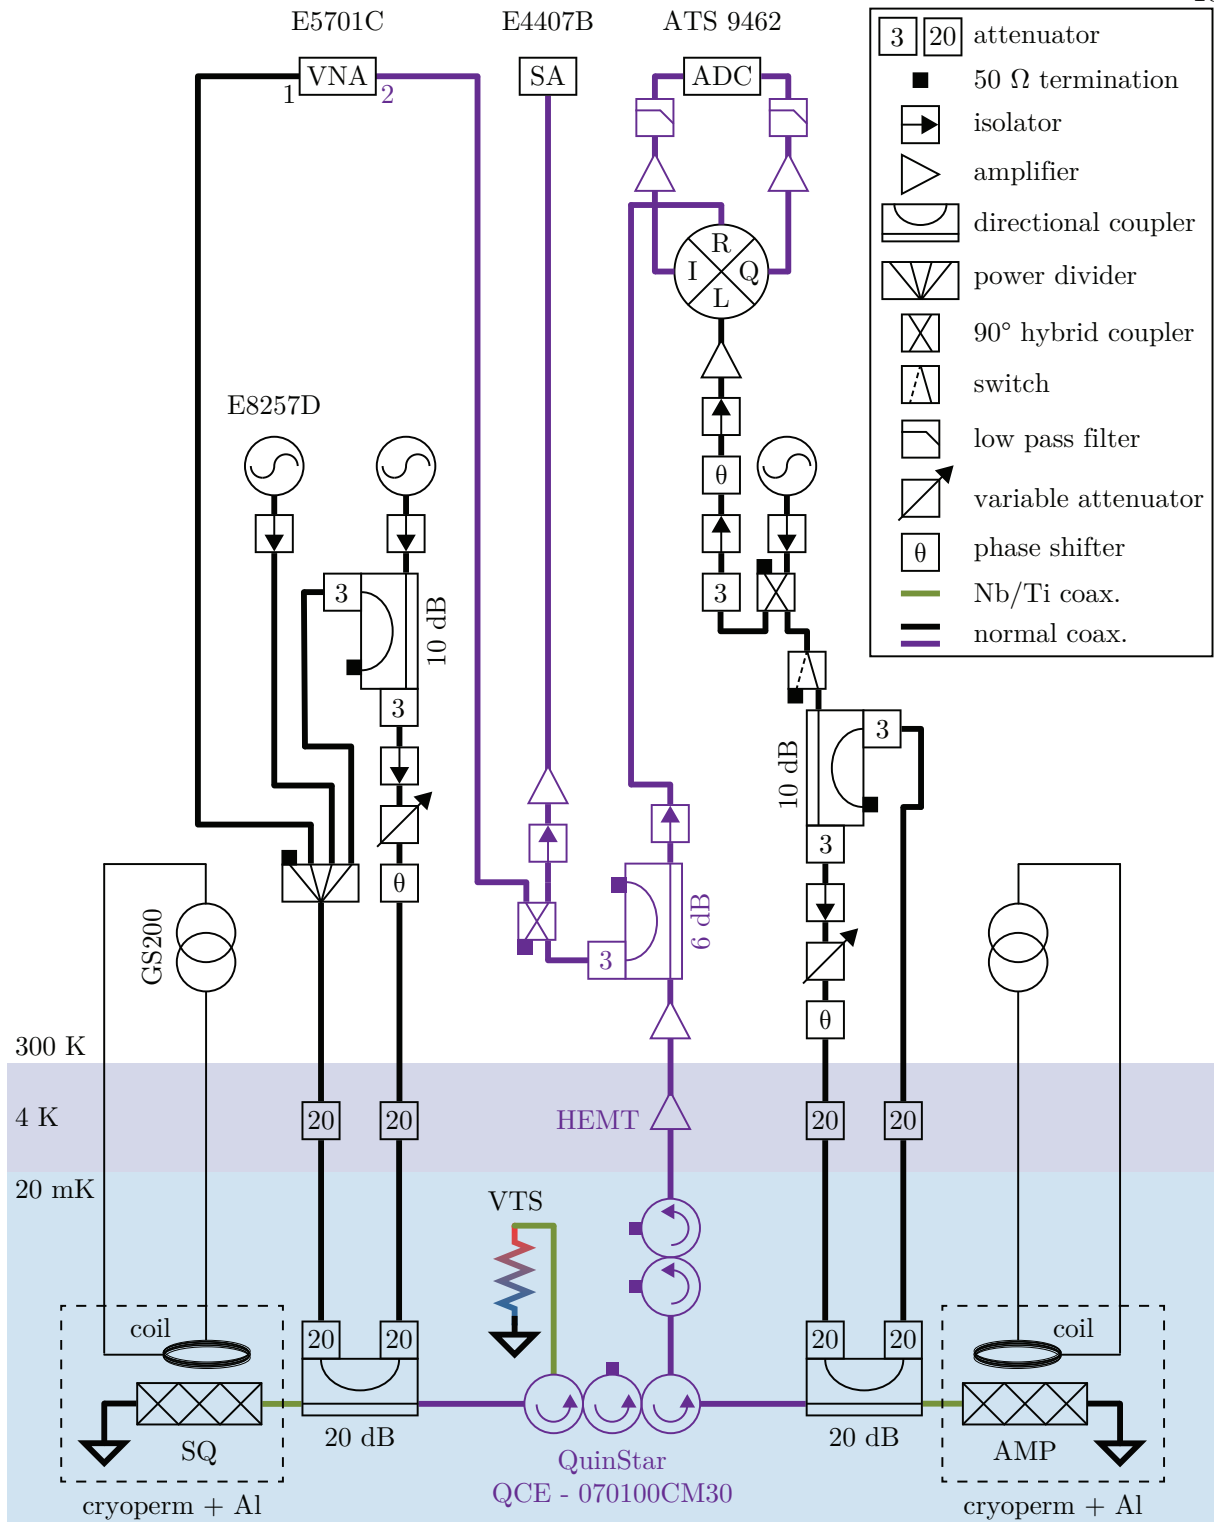


Figure 4.7: Full experimental setup used for both coherent tone deamplification and squeezing. The input and output coaxial cables are colored in black and purple, respectively. The green lines are superconducting. Figure reproduced from Ref. [204].

4.4.2 Inter-JPA loss calibration

Calibrating the lines relies on measuring the overall input attenuation A^I from the probe tone's microwave generator to the SQ input, and also the overall (single-quadrature) amplifications $G_{A,1Q}^O$ and $G_{S,1Q}^O$ from the TKC outputs to the analog-to-digital converter (ADC). Knowing A^I allows for calibration of the probe power incident on the SQ TKC, and the ratio

$$\eta = G_{A,1Q}^O / G_{S,1Q}^O \quad (4.5)$$

provides, *ceteris paribus*, an estimate of the loss experienced by the squeezed state when travelling between SQ and AMP.

Operating one JPA at a time (without pump cancellation), $G_{A,1Q}^O$ and $G_{S,1Q}^O$ are calculated by integrating the power spectral density of vacuum fluctuations recorded by the ADC over a window $W = 500$ kHz within the JPA's bandwidth. A heterodyne measurement, with LO and JPA pump detuned by 5 MHz, and W ending 400 kHz below the pump, is performed to avoid DC offsets. Thus, when measuring $G_{A,1Q}^O$, the generator used for pumping is the one connected to the SQ, itself detuned. The SQ TKC then simply acts as a mirror, and the pump tone reaches the AMP after reflection. Given a quantum at the SQ or AMP input, the integrated power for both cases is

$$P_{S/A}^O = \hbar\omega W G_{S/A}^O G_{S/A}, \quad (4.6)$$

with $G_{S/A}$ the SQ or AMP gain, measured with the VNA. Operating the JPAs at various gains between 20 and 30 dB in order to ensure that the output spectral density linearly tracked the JPA gain yielded

$$G_{S,1Q}^O = 75.3 \pm 0.1 \text{ dB} \quad (4.7)$$

$$G_{A,1Q}^O = 76.5 \pm 0.1 \text{ dB}. \quad (4.8)$$

Therefore, we estimate that $\eta = 1.2 \pm 0.2$ dB. In principle this value should enable us to observe about 5.6 dB of squeezing, a bit more than what was measured. But in practice, we are limited by a combination of squeezed state distortion and the contribution from the high-electron-mobility

transistor (HEMT) amplifier added noise, which becomes harder to overwhelm when squeezing because the HEMT added noise is larger relative to the squeezed state than to vacuum.

To estimate A^I , we use a similar protocol. Turning off both JPAs, a probe tone with input power P_P^I travelling in the lines creates an output power P_P^O on the ADC as

$$P_P^O = P_P^I G_{S,1Q}^O A^I, \quad (4.9)$$

leading to an estimate of $A^I = -81.4 \pm 0.2$ dB.

4.4.3 Added noise calibration

Any resistor at temperature T generates noise of known variance, whose spectral density in units of quanta at angular frequency ω is

$$S_{\text{in}} = \frac{1}{2} + \frac{1}{e^{\hbar\omega/k_B T} - 1}. \quad (4.10)$$

The VTS therefore allows us to generate a known thermal state at the chain's input, which is then amplified by the SQ, giving at its output:

$$S_{S,\text{out}} = G_{S,1Q}(S_{\text{in}} + N_S), \quad (4.11)$$

where $G_{S,1Q}$ is the SQ gain and N_S its added noise (in this experiment only the SQ is used; the AMP is detuned). Since we are operating the SQ in a phase-sensitive mode, $N_S \geq 0$.

In practice there are inefficiencies along the thermal state's path from VTS to SQ. We can model these with a simple beamsplitter picture, such that a fraction λ of S_{in} is replaced with noise S_f at the fridge's temperature. Considering the system added noise as

$$N_{\text{add}} = N_S + \frac{N_H}{G_{S,1Q}}, \quad (4.12)$$

where N_H is the HEMT added noise, we have

$$\begin{aligned} S_{\text{out}} &= G_S^O G_S (\lambda S_{\text{in}} + (1 - \lambda) S_f + N_{\text{add}}) \\ &= G_S^O G_S S_{\text{out}}^{\text{in}} \end{aligned} \quad (4.13)$$




Figure 4.8: Noise spectral density as a function of variable temperature stage temperature. The noise is referred to the SQ input. Points in black, blue, and red are for fridge’s temperatures $T_f = 50, 300,$ and 500 mK, respectively. The statistical error is smaller than the point size. The solid lines are obtained from fitting the parameters N_{add}, λ and the overall gain of the chain $G_{S,1Q}^O G_S$. Figure reproduced from Ref. [204].

at the output of the chain, where $G_{S,1Q}^O$ is the chain’s gain, and $S_{\text{out}}^{\text{in}}$ is the output noise spectral density referred to the SQ input. Varying S_{in} and S_f , we thus can deduce N_{add} and λ .

With a homodyne configuration (LO and SQ pump at the same frequency), we integrate the output noise spectral density S_{out} over the SQ bandwidth B for various VTS temperatures. The fridge’s base temperature T_f is fixed and $G_{S,1Q} \geq 25$ dB is tracked by amplifying and measuring a small pilot tone. Figure 4.8 shows $S_{\text{out}}^{\text{in}}$ as a function of the VTS temperature for three different fridge temperatures. A fit gives $N_{\text{add}} = 0.045 \pm 0.001$ and $\lambda = 0.79 \pm 0.01$, where the quoted uncertainties correspond to statistical errors in the measurements.

4.5 Concluding remarks on current-pumped JPAs optimized for squeezing

Pump stiffness is crucial when generating a microwave squeezed state with a JPA, because any deviation from the stiff-pump regime strongly affects the squeezed quadrature. A lack of

stiffness distorts output signals, thereby limiting the maximum degree of squeezing. We have seen that operating above the LMG therefore constitutes a partial solution to the problem of finite pump stiffness by means of choosing the operating point to lie above the LMG. The coherent tone deamplification measurements of Sec. 4.2 confirm the efficacy of this approach, with less phasor distortion above the LMG. When squeezing vacuum, we see the same trend. Maintaining low microwave loss, it was thus possible to generate and deliver a 3.87 ± 0.03 dB vacuum squeezed state from one JPA to another over a band of 30 MHz.⁹¹

The benefits of optimizing JPA operation explored in this chapter will grow as losses shrink. With relatively high losses, it does little good to create a minimally deformed squeezed state, only to have most of it replaced by unsqueezed vacuum. However, as losses come down for experimentalists, preventing squeezed state deformation will have impact beyond the few percent level achieved in this chapter, as observed in Sec. 4.2 for the deamplification of a coherent tone.

For the purposes of axion detection with present-day technology, however, the gains achieved in this chapter are insufficient to recommend current-pumped JPAs in light of their two key disadvantages for axion detection. First, the losses introduced by the directional couplers in Fig. 4.7 degrade the itinerant squeezed state more than is absolutely necessary.⁹² Second, the pump cancellation of Sec. 4.3 poses a significant experimental challenge for an active haloscope which must be automatically tuned through frequency space. This difficulty is further compounded by the Kerr shift of current-pumped JPAs. Unlike with flux-pumped JPAs, changing the drive power changes the resonant frequency, making it harder still to stabilize both pump and cancellation tones for two JPAs simultaneously. In light of this conclusion, subsequent work in this thesis uses the flux-pumped JPAs of Sec. 3.7.

⁹¹ This 30 MHz is more than enough for axion haloscope purposes: as we will see in Chs. 5 and 8, the usable bandwidth for an axion haloscope at these frequencies is closer to 1–2 MHz.

⁹² While switching to flux-pumped JPAs will get rid of these directional couplers, the introduction of the second, flux-pump port on the TKCs does allow for signal loss. In practice, however, the amount of loss out this port (Fig. 3.27b) is far less than that from the reflection off and transmission through a directional coupler.

Chapter 5

Squeezed vacuum used to accelerate the search for a weak classical tone

In Ch. 2, we saw that the present generation of axion haloscope experiments [91, 107, 108] is or soon will be limited by vacuum fluctuations in the microwave field. The quadrature operators \hat{X} and \hat{Y} of a resonant cavity's electromagnetic (EM) field, Eqs. (3.137) and Eqs. (3.138), carry the imprint of the dark matter signal as a slight excess in their power spectra. Problematically, the quantum noise [128] intrinsic to a measurement of these observables overpowers the signal by roughly three orders of magnitude,⁹³ and the signal frequency is *a priori* unknown [226]. Among the most promising advances are those which allow access to a fundamentally distinct parameter regime in which axion searches are no longer limited by quantum noise.

In detectors measuring the quadratures of a resonant mode, quantum noise can be circumvented by preparing the mode in a squeezed state. Squeezing unbalances the uncertainties in the two quadrature observables, thereby permitting precise knowledge of one at the expense of the other. Leveraging quantum squeezing to enhance measurement sensitivity has been a longstanding goal [227–230] in the optical domain, now realized for the sensing of gravitational waves [231]. At microwave frequencies, squeezing had, prior to the work described in Ch. 8 of this thesis, been demonstrated in principle [209, 213, 214], without yet aiding in a search for new physical phenomena.

Microwave squeezing can be used to aid in an axion dark matter search subject to current experimental constraints. This chapter presents the work of Ref. [208] in order to substantiate that

⁹³ A three order-of-magnitude discrepancy between signal and noise spectral densities requires six orders of magnitude worth of independent realizations of the axion field merely to achieve an SNR of 1 [77], itself insufficient to detect or exclude the axion.

claim, first theoretically and then experimentally, before the conclusive experiment is performed in Ch. 8. In particular, we will see in this chapter that squeezing increases the scan rate in a search for a weak, axion-like signal of unknown frequency. This improvement exists in spite of the fact that squeezing does not improve the sensitivity of a haloscope to a tone of known frequency. This speed-up is demonstrated in an apparatus designed to mimic the behavior of existing haloscopes. With a maximum of 4.5 ± 0.1 dB of squeezing delivered from one flux-pumped Josephson parametric amplifier (JPA; Sec. 3.7) to another in a squeezed state receiver (SSR) configuration, we will see implemented a realistic acquisition and processing protocol both with and without squeezing enabled. The SSR presented in this chapter accelerates the scan rate to axion-like signals by a factor of 2.12 ± 0.08 over 200 independent experimental trials. In Ch. 8, we will equip the HAYSTAC experiment with the SSR developed in this chapter to speed it up by very nearly that amount.

5.1 Theory of the squeezed state receiver

Figure 5.1 shows a representative experimental apparatus in which a resonant cavity is coupled to an SSR comprising a pair of JPAs. We start by considering the cavity, whose internal mode (frequency ω_c) has Hamiltonian

$$\hat{H} = \hbar\omega_c \frac{\hat{X}^2 + \hat{Y}^2}{2}, \quad (5.1)$$

equivalent to that of Eq. (3.62) subject to the definitions of Eqs. (3.137) and (3.138). The quadratures of the cavity field obey

$$[\hat{X}, \hat{Y}] = i, \quad (5.2)$$

a direct consequence of their definitions and Eq. (3.59). This cavity is modeled as exchanging energy with three ports. First, a measurement port couples the cavity mode to the propagating modes of a transmission line with power decay rate κ_m . Along this line, a microwave circulator spatially separates incoming and outgoing propagating modes. Second, a loss port, connected to a fictitious transmission line, models the cavity's internal energy dissipation at rate κ_l . Third, the cavity's coupling to the signal of interest at rate κ_a is modeled as occurring through another

fictitious transmission line; the signal itself is modeled as a microwave generator characterized by its frequency ω_a , spectral width⁹⁴ Δ_a , and amplitude \mathcal{E}_a . We assume $\mathcal{E}_a \gg 1$, implying that the displacement of the cavity mode by the signal is classical (i.e., the contribution of the signal to the cavity's quantum fluctuations can be neglected in comparison to those coming from the loss and measurement ports). We also assume a narrowband signal ($\Delta_a \ll \{\kappa_l, \kappa_m\}$) so weakly coupled ($\mathcal{E}_a^2 \kappa_a \ll \kappa_l$) that the time required to resolve the displacement is much longer than the signal's phase coherence time. Therefore, on average, this displacement yields a small excess power above the vacuum fluctuations, isotropic in quadrature space (bottom panel of Fig. 5.1). These inequalities are very well satisfied in the case of the axion field (Sec. 5.2).

The SSR itself comprises two JPAs coupled respectively to incoming and outgoing modes at the cavity's measurement port. This configuration exploits the fact that a portion of the vacuum noise exiting the measurement port arises from vacuum noise incident on that same port. The first JPA (called SQ, as in Ch. 4) squeezes these input fluctuations along the \hat{X} quadrature, reducing the observable's variance below vacuum levels:

$$\sigma_X^2 < \frac{1}{2}. \quad (5.3)$$

To satisfy the uncertainty principle, the opposing quadrature's fluctuations exceed vacuum

$$\sigma_Y^2 > \frac{1}{2}, \quad (5.4)$$

but are not measured. The squeezed input field subsequently enters the cavity, where a small displacement by the signal would yield a small excess power in both quadratures. At the measurement port output, the second JPA (again called AMP) noiselessly amplifies only the \hat{X} quadrature with sufficient gain to overwhelm the noise added by following amplifiers and mixers.

The benefit of squeezing can be understood by analyzing the microwave network formed by the combination of SSR and cavity using input-output theory (Sec. 3.4.4). We will here qualitatively overview the results of more detailed calculations that we will perform first in the lossless case in

⁹⁴ A closely related quantity is the standard deviation of the axion lineshape, σ_{SHM} in Fig. 2.5. The numbers used the quote the axion spectral width traditionally encompass a bit more of the axion lineshape than the standard deviation does.




Figure 5.1: Top panel: schematic of the SSR and cavity. Two Josephson parametric amplifiers, squeezer (SQ) and amplifier (AMP), respectively squeeze and read out a microwave field interacting with a cavity at rate κ_m . An axion-like field \mathcal{E}_a , coupled to the cavity at rate κ_a , displaces the cavity state. Energy leaves the cavity through internal absorption at a rate κ_l . Bottom panel: quadrature representation of a vacuum state detuned from the cavity resonant frequency, travelling through the SSR. At the SQ’s input (a) it is Gaussian, azimuthally equiprobable in the (\hat{X}, \hat{Y}) plane (red disk). The state is squeezed along \hat{X} by the SQ (b), displaced along a random phase within the cavity (c), and amplified along \hat{X} by the AMP (d). Comparing to what happens without squeezing (e, f and g), the size of the signal-plus-noise (green) relative to the noise (red) in this quadrature is larger with (d) than without (g) squeezing. Figure reproduced from Ref. [208].

Sec. 5.1.1, and then with experimental transmission losses modelled in Sec. 5.1.2. These calculations show the benefit of squeezing against unsqueezed, single-quadrature axion detection. The picture is completed in Sec. 5.1.3, where we will see that, in the absence of squeezing, there is neither a benefit nor a penalty associated with measuring only one quadrature compared to the usual two-quadrature case. Before describing experimental results in the rest of the chapter, in Sec. 5.2 we will map parameters of the model shown in Fig. 5.1 onto the physical parameters of e.g. Eq. (2.16), which are conventionally quoted when searching for dark matter axions.

The behavior of the apparatus shown in Fig. 5.1 is characterized by its signal and noise spectral densities. The signal spectral density at the measurement port output is equal to the signal spectral density at the measurement port input, weighted by the susceptibility of the measurement port output to the signal port input. Similarly, the noise density at the measurement port output is a susceptibility-weighted sum of squeezed and unsqueezed noise from the measurement and loss ports, respectively. In the absence of transmission losses between the JPAs and the cavity, we will see that the ratio of output signal spectral density to total output noise spectral density (hereafter called the signal visibility) is

$$\alpha(\omega) \approx \frac{n_A \kappa_a \kappa_m}{(n_T + \frac{1}{2}) \left(\kappa_l \kappa_m + \frac{\beta(\omega)}{G_{S,1Q}} \right)}. \quad (5.5)$$

Here, ω is the frequency relative to cavity resonance,

$$n_T = \frac{1}{e^{\hbar\omega_c/k_B T} - 1} \quad (5.6)$$

is the mean thermal photon number incident on the cavity from the measurement and loss ports, $n_A = \mathcal{E}_a^2$ is the mean photon number sourced by the fictitious generator,

$$\beta(\omega) = \left(\frac{\kappa_m - \kappa_l}{2} \right)^2 + \omega^2, \quad (5.7)$$

and $G_{S,1Q}$ is the single-quadrature power gain of the SQ, Eq. (3.139), along its amplified quadrature (ideally equal to the reduction of the squeezed state variance below the vacuum value).

When optimizing $\alpha(\omega)$ in Eq. (5.5), we assume that κ_m and $G_{S,1Q}$ can be freely varied, whereas n_T , n_A , κ_l , and κ_a are fixed by the physics of the signal source or technical constraints of the detector. On cavity resonance ($\omega = 0$), α is maximized at critical coupling ($\kappa_m = \kappa_l$). Because $\beta(0) = 0$ at critical coupling, $\alpha(0)$ is independent of $G_{S,1Q}$ and there is no benefit from squeezing; physically, the squeezed state injected into the cavity is completely absorbed in it, while all the unsqueezed noise from the loss port reaches the AMP. For $\omega \neq 0$, squeezing increases $\alpha(\omega)$ for any value of κ_m , as $G_{S,1Q}$ reduces the amount of measurement port noise reaching the output. In the limit where $G_{S,1Q} \rightarrow \infty$, the $\beta(\omega)$ term in the denominator can be neglected, and $\alpha(\omega)$ approaches

the critically coupled resonant value $\alpha(\omega = 0, \kappa_m = \kappa_l)$ for *any* κ_m and *all* ω . This illustrates an important point of principle: squeezing cannot improve the peak sensitivity of a haloscope, but it can almost without limit⁹⁵ enhance the detector bandwidth over which this peak sensitivity is achieved. When $G_{S,1Q}$ is finite, overcoupling ($\kappa_m > \kappa_l$) increases the cavity bandwidth at the cost of reducing $\alpha(0)$. This can be a favorable tradeoff (even in the absence of squeezing) because the signal's frequency ω_a is *a priori* unknown [60, 95], and broader bandwidth enables larger cavity tuning steps. Moreover, squeezing mitigates the reduction of $\alpha(0)$ from overcoupling, thus enabling faster tuning without significant degradation of sensitivity.

To quantify this speed-up, we will also calculate the scan rate R , Eq. (2.16), at which we can tune the cavity resonance through frequency space in search for a signal. This scan rate is inversely proportional to the measurement time at each tuning step, which in turn scales with α^{-2} as a consequence of Gaussian noise statistics (Sec. 5.1.1). Hence

$$R \propto \Delta_a \int_{-\infty}^{\infty} \alpha^2(\omega) d\omega; \quad (5.8)$$

carrying out the integral we obtain

$$R \propto \frac{\Delta_a \sqrt{G_{S,1Q}} n_A^2 \kappa_a^2 \kappa_m^2}{(n_T + \frac{1}{2})^2 \left[\kappa_l \kappa_m + \frac{1}{G_{S,1Q}} \left(\frac{\kappa_l - \kappa_m}{2} \right)^2 \right]^{3/2}}. \quad (5.9)$$

Without squeezing ($G_{S,1Q} = 1$), R is maximized when the cavity is twice-overcoupled ($\kappa_m = 2\kappa_l$), and at this optimal coupling the scan rate scales as $R_u^{\max} \propto \kappa_a^2 / \kappa_l$, a known result [90].⁹⁶ When $G_{S,1Q} \gg 1$, the optimal coupling is $\kappa_m = 2G_{S,1Q}\kappa_l$ and the scan rate scales as $R_s^{\max} \propto G_{S,1Q}\kappa_a^2 / \kappa_l$. Comparing the two situations, the scan rate is improved by $G_{S,1Q}$, which shows that an ideal SSR greatly accelerates the search for a weak classical signal when squeezing and overcoupling.

In practice, however, we will see in Sec. 5.1.2 how losses in microwave components reduce the SQ-to-AMP transmission efficiency η and hence the benefit of squeezing, because part of the

⁹⁵ Presumably, once the receiver bandwidth approaches the frequency of the cavity, it would become inviable to measure the cavity state faster than its own oscillations. However, this limit is over three orders of magnitude away from the performance achieved here and in Ch. 8.

⁹⁶ In particular, a twice-overcoupled haloscope (without squeezing) scans $32/27 \approx 1.19\times$ faster than the same haloscope in its critically coupled configuration, despite losing a little sensitivity at the cavity resonance frequency.

squeezed state is replaced with unsqueezed vacuum. Figure 5.2 compares the theoretical scan rate enhancement $E_t = R_s/R_u^{\max}$ (a) when $\eta = 1$ (perfect transmission), and (b) when $\eta = 0.69$ (efficiency observed in practice), as a function of $G_{S,1Q}$ and κ_m/κ_l . In the first case, E_t improves arbitrarily as squeezing and coupling are together increased. In the second case, it plateaus at $E_t^{\max} \approx 2.2$ for $G_{S,1Q} > 20$ when optimally coupled. In Secs. 5.3 and 5.4, we will see experimental results for the scan rate enhancement consistent with this theory.




Figure 5.2: (a) Theoretical scan rate enhancement E_t , calculated as a function of SQ single-quadrature power gain $G_{S,1Q}$ and coupling ratio κ_m/κ_l . With perfect efficiency, i.e. $\eta = 1$, scan rate improves steadily with $G_{S,1Q}$ and κ_m . (b) Same as (a) for $\eta = 0.69$. Scan rate plateaus at 2.2 for $\kappa_m/\kappa_l = 5$ and $G_{S,1Q} > 20$. The color scales for (a) and (b) differ by close to a factor of 10. Figure adapted from Ref. [208] to correct an error whereby the lowest value on the color axes were labeled “1.” The other axis labels have been verified.

5.1.1 Lossless calculation

In this section, we track the propagating EM fields through the SSR and cavity, so as to derive the susceptibility matrix of the entire system. We then calculate the signal visibility $\alpha(\omega)$ and the scan rate R . We first consider the case where the propagating fields experience no loss (the cavity mode still decays partially out its loss port), which nonetheless faithfully illustrates the utility of the SSR, then treat the full system in the presence of transmission losses in Sec. 5.1.2.

We model the energy exchange between the cavity’s ports in its own rotating frame. The

time evolution of the cavity is governed by the Heisenberg-Langevin equation

$$\frac{d\hat{A}}{dt} = -\frac{\kappa_T}{2}\hat{A}(t) + \sum_j \sqrt{\kappa_j}\hat{a}_{\text{in},j}(t), \quad (5.10)$$

where \hat{A} is the cavity ladder operator,

$$\kappa_T = \kappa_m + \kappa_l + \kappa_a, \quad (5.11)$$

and $\hat{a}_{\text{in},j}$ ($j = m, l, a$) are the annihilation field operators of the input modes incident on ports indexed by m, l , and a . At the measurement port, we physically separate input ($\hat{a}_{\text{in},m}$) and output ($\hat{a}_{\text{out},m}$) fields with a circulator. Given the input-output (or boundary conditions) relations

$$\sqrt{\kappa_j}\hat{A}(t) = \hat{a}_{\text{in},j}(t) - \hat{a}_{\text{out},j}(t), \quad (5.12)$$

analogous to those of Eq. (3.100), the input and output field operators are related in the Fourier domain by

$$\hat{a}_{\text{out},j}(\omega) = \sum_k \chi_{jk}(\omega)\hat{a}_{\text{in},k}(\omega), \quad (5.13)$$

where

$$\chi_{jk}(\omega) = \frac{-\sqrt{\kappa_j\kappa_k} + (\kappa_T/2 + i\omega)\delta_{jk}}{\kappa_T/2 + i\omega} \quad (5.14)$$

are the elements of a 3×3 susceptibility matrix fully describing the behavior of the cavity [200].

By cascading the input-output relations for each element presented in Fig. 5.1, we can calculate the benefit from squeezing. We work in the quadrature basis and consider the vector of input quadratures

$$\vec{x}_{\text{in}} = [\hat{X}_{\text{in},m}, \hat{X}_{\text{in},l}, \hat{X}_{\text{in},a}]^T \quad (5.15)$$

aligned with our squeezing. We will calculate the SSR/cavity susceptibility matrix $\Xi_{\mathbf{X}}$, in terms of which the vector of output quadratures is

$$\vec{x}_{\text{out}} = \Xi_{\mathbf{X}}\vec{x}_{\text{in}}. \quad (5.16)$$

The first element in the system is the SQ, which performs a one-mode squeezing (OMS) operation on the measurement port's input quadrature:

$$\vec{x}_S = \mathbf{S}_X \vec{x}_{\text{in}} = \begin{bmatrix} \frac{1}{\sqrt{G_{S,1Q}}} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \vec{x}_{\text{in}}, \quad (5.17)$$

where the subscript S refers to the SQ output port and $G_{S,1Q}$ is the SQ single-quadrature power gain. The OMS operation also amplifies the other quadrature of the measurement port mode, $\hat{Y}_{\text{in},m}$ by $\sqrt{G_{S,1Q}}$ in order to preserve the Heisenberg uncertainty relation. We do not track the evolution of the \hat{Y} quadrature here, as it is irrelevant for SSR performance.

Next, the cavity transforms the quadrature operators. In the cavity rotating frame, the vectors of quadrature operators are obtained from the vectors of ladder operators by

$$\begin{bmatrix} \vec{x} \\ \vec{y} \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} \mathbf{I}_3 & \mathbf{I}_3 \\ -i\mathbf{I}_3 & i\mathbf{I}_3 \end{bmatrix} \begin{bmatrix} \vec{a}(\omega) \\ \vec{a}^\dagger(-\omega) \end{bmatrix}, \quad (5.18)$$

or

$$[\vec{x}, \vec{y}]^\text{T} = \mathbf{P}[\vec{a}(\omega), \vec{a}^\dagger(-\omega)]^\text{T} \quad (5.19)$$

Here \mathbf{I}_3 is the 3×3 identity matrix, and the mode ordering in each vector is m, l, a : for example,

$$\vec{a}(\omega) = [\hat{a}_m(\omega), \hat{a}_l(\omega), \hat{a}_a(\omega)]^\text{T}. \quad (5.20)$$

Thus, the cavity susceptibility matrix in the quadrature basis $\tilde{\chi}$ is:

$$\tilde{\chi} = \mathbf{P} \begin{bmatrix} \chi(\omega) & \mathbf{0}_3 \\ \mathbf{0}_3 & \chi^*(-\omega) \end{bmatrix} \mathbf{P}^{-1} = \begin{bmatrix} \chi(\omega) & \mathbf{0}_3 \\ \mathbf{0}_3 & \chi(\omega) \end{bmatrix}, \quad (5.21)$$

where $\mathbf{0}_3$ is the null matrix. We see that the cavity's effect on field and quadrature operators is identical, which leads to

$$\vec{x}_o = \chi(\omega) \vec{x}_S, \quad (5.22)$$

where o refers to the cavity output port.

Finally the AMP performs a second OMS operation so as to amplify the quadrature that the SQ originally squeezed:

$$\vec{x}_{\text{out}} = \mathbf{A}_{\mathbf{X}} \vec{x}_o = \begin{bmatrix} \sqrt{G_{A,1Q}} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \vec{x}_o, \quad (5.23)$$

where $G_{A,1Q}$ is the AMP power gain. The SSR and cavity thus transform input to output quadratures according to

$$\vec{x}_{\text{out}} = \mathbf{A}_{\mathbf{X}} \chi(\omega) \mathbf{S}_{\mathbf{X}} \vec{x}_{\text{in}} = \mathbf{\Xi}_{\mathbf{X}} \vec{x}_{\text{in}}. \quad (5.24)$$

We can then calculate the single-quadrature output spectral density matrix

$$\mathbf{\Sigma}_{\text{out},\mathbf{X}} = \frac{\langle [\vec{x}_{\text{out}}]^\dagger [\vec{x}_{\text{out}}]^T \rangle}{2\pi}, \quad (5.25)$$

where the Hermitian conjugation here does not transpose the vector, in order to determine the total output signal and noise powers. Substituting for \vec{x}_{out} yields

$$\mathbf{\Sigma}_{\text{out},\mathbf{X}} = \mathbf{\Xi}_{\mathbf{X}}^* \mathbf{\Sigma}_{\text{in},\mathbf{X}} \mathbf{\Xi}_{\mathbf{X}}^T, \quad (5.26)$$

where

$$\mathbf{\Sigma}_{\text{in},\mathbf{X}} = \frac{\langle [\vec{x}_{\text{in}}]^\dagger [\vec{x}_{\text{in}}]^T \rangle}{2\pi} \quad (5.27)$$

is the input noise spectral density matrix. Since the SSR is connected to three incoherent and uncorrelated modes (Fig. 5.1),

$$\frac{\langle \hat{X}_{\text{in},j} \hat{X}_{\text{in},k} \rangle}{2\pi} = \delta_{jk} \left(n_{\text{in},j} + \frac{1}{2} \right), \quad (5.28)$$

where $n_{\text{in},j}$ is the input mean photon number per unit time per unit bandwidth (henceforth simply mean photon number) of mode j . Thus,

$$\mathbf{\Sigma}_{\text{in},\mathbf{X}} = \begin{bmatrix} n_T + \frac{1}{2} & 0 & 0 \\ 0 & n_T + \frac{1}{2} & 0 \\ 0 & 0 & n_A + \frac{1}{2} \end{bmatrix}. \quad (5.29)$$

Taking the first entry in the matrix $\Sigma_{\text{out},\mathbf{X}}$ we obtain the single-quadrature output spectral density at the measurement port:

$$\Sigma_{\text{out},X,m} \approx \frac{G_{A,1Q}}{B(\omega)} \left[n_A \kappa_a \kappa_m + \left(n_T + \frac{1}{2} \right) \left(\kappa_l \kappa_m + \frac{\beta(\omega)}{G_{S,1Q}} \right) \right], \quad (5.30)$$

where

$$B(\omega) = \left(\frac{\kappa_m + \kappa_l}{2} \right)^2 + \omega^2, \quad (5.31)$$

$\beta(\omega)$ is given by Eq. (5.7), and we used $\kappa_a \ll \{\kappa_l, \kappa_m\}$ and $n_A \gg 1/2$.

The visibility $\alpha(\omega)$ can be directly extracted from $\Sigma_{\text{out},X,m}$, since the signal is the term proportional to n_A and the noise is $\Sigma_{\text{out},X,m}(n_A = 0)$. The result is exactly that quoted in Eq. (5.5).

We now consider a search protocol comprising many measurements of $\Sigma_{\text{out},X,m}$, as the cavity frequency is tuned by discrete steps δ_c . The spectral scan rate R is obtained by taking the limit:

$$R = \lim_{\substack{\delta_c \rightarrow 0 \\ \tau \rightarrow 0}} \frac{\delta_c}{\tau}, \quad (5.32)$$

where τ is the duration of each measurement. In the rest of this section we will derive Eq. (5.9) for the scan rate. We begin by defining the signal-to-noise ratio (SNR) at a single tuning step as

$$\bar{\alpha}(\omega) = \sqrt{\tau \Delta_a} \frac{\alpha(\omega)}{2}, \quad (5.33)$$

which scales as the square root of the number of independent measurements of the power contained in the bandwidth Δ_a , centered on a frequency detuned from cavity resonance by $\omega/2\pi$ [77, 106]. The factor of 2 is a consequence of our single-quadrature measurement scheme, which is discussed in Sec. 5.1.3. The SNR is a useful quantity, as it relates the visibility to how much time is actually needed to resolve an axion signal.

For $\delta_c \lesssim \kappa_T/2\pi$, multiple tuning steps will contribute to the SNR at each frequency. Without loss of generality, we evaluate the net SNR for some putative signal frequency ω_a which coincides exactly with the cavity resonance at a particular step. We then consider the contributions to the SNR of the surrounding $2n + 1$ tuning steps, where

$$2\pi\delta_c = 2\kappa_T/(2n + 1). \quad (5.34)$$

Contributions to the SNR add in quadrature, so the net squared SNR is

$$\bar{\alpha}_{2n+1}^2 = \frac{\tau \Delta_a}{4} \frac{1}{\delta_c} \sum_{k=-n}^{k=n} \alpha^2(k2\pi\delta_c) \delta_c. \quad (5.35)$$

In the limit $n \rightarrow \infty$, $\tau \rightarrow 0$, we obtain the integrated squared SNR

$$\bar{\alpha}_I^2 = \frac{\Delta_a}{4} \frac{1}{R} \int_{-\infty}^{\infty} \alpha^2(\omega) \frac{d\omega}{2\pi}, \quad (5.36)$$

where we have extended the limits of integration to $\pm\infty$: the contributions from tuning steps where the cavity is detuned by more than κ_T are negligible.

Note that our arbitrary choice of ω_a appears nowhere explicitly on the right-hand side of Eq. (5.36); thus the integrated SNR will be frequency-independent, insofar as the quantities contributing to $\alpha(\omega)$ remain constant as we tune the cavity. In practice, we set a certain target value of $\bar{\alpha}_I$ for which a real signal would appear as a sufficiently prominent peak in the grand spectrum (see e.g., Fig. 5.7b); this choice determines the required spectral scan rate R . Solving for R , we obtain

$$\begin{aligned} R &= \frac{\Delta_a}{4\bar{\alpha}_I^2} \int_{-\infty}^{\infty} \alpha^2(\omega) \frac{d\omega}{2\pi} \\ &= \frac{\Delta_a \sqrt{G_{S,1Q}} n_A^2 \kappa_a^2 \kappa_m^2}{16\bar{\alpha}_I^2 \left(n_T + \frac{1}{2}\right)^2 \left[\kappa_l \kappa_m + \frac{1}{G_{S,1Q}} \left(\frac{\kappa_l - \kappa_m}{2}\right)^2\right]^{3/2}}, \end{aligned} \quad (5.37)$$

in agreement with Eq. (5.9). In a comparison of squeezed and unsqueezed scan rates, the $\Delta_a/4\bar{\alpha}_I^2$ factor drops out; thus, we need only compare the integral in Eq. (5.9) for the two cases.

5.1.2 Calculation with transmission losses

Here we generalize the calculations of Sec. 5.1.1 to account for imperfect power transmission efficiency η between SQ and AMP. Integrating the squared SNR over all frequencies, we obtain the scan rate as a function of η , used in Fig. 5.2. For simplicity, we treat the transmission efficiency λ between the SQ and the cavity as identical to that between the cavity and the AMP, hence $\eta = \lambda^2$. Note that in principle, loss between SQ and cavity is slightly less harmful than loss between cavity and AMP, as loss that occurs after the cavity degrades the signal along with the squeezing.

As in Sec. 5.1.1, we track the vector of input quadratures \vec{x}_{in} throughout the SSR and cavity system, neglecting the orthogonal quadratures. First, the SQ performs the OMS operation defined by Eq. (5.17): $\vec{x}_S = \mathbf{S}_X \vec{x}_{\text{in}}$. Next we must account for transmission losses between the SQ and the cavity, modeled as a beamsplitter interaction between the quadrature operators of the measurement line at the SQ output $\{\hat{X}_{s,m}, \hat{Y}_{s,m}\}$, and those of another, uncontrolled mode $\{\hat{X}_{s,\lambda}, \hat{Y}_{s,\lambda}\}$. We could define generalized quadrature vectors \vec{x}_S and \vec{y}_S to include $\hat{X}_{s,\lambda}$ and $\hat{Y}_{s,\lambda}$ respectively, but we do not need to keep track of the EM field's evolution in loss modes. Among the modes we do keep track of, only the mode entering and exiting the measurement port experiences the loss:

$$\vec{x}_i = \begin{bmatrix} \sqrt{\lambda} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \vec{x}_S, \quad (5.38)$$

or $\vec{x}_i = \mathbf{L}_X \vec{x}_S$. Here i refers to the cavity's input, and $0 \leq \lambda \leq 1$ is the single-sided transmission efficiency.

Since we are not keeping explicit track of the propagating loss mode between the SQ and cavity in this calculation, we account for the unsqueezed vacuum that it introduces as an added noise term,

$$\mathbf{N}_X = \begin{bmatrix} (n_T + \frac{1}{2})(1 - \lambda) & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \quad (5.39)$$

We thus have the single-quadrature noise spectral density at the cavity input:

$$\boldsymbol{\Sigma}_{i,X} = [\mathbf{L}_X \mathbf{S}_X]^* \boldsymbol{\Sigma}_{\text{in},X} [\mathbf{L}_X \mathbf{S}_X]^T + \mathbf{N}_X, \quad (5.40)$$

where $\boldsymbol{\Sigma}_{\text{in},X}$ is given by Eq. (5.29). This expression yields

$$\boldsymbol{\Sigma}_{i,X} = \begin{bmatrix} (n_T + \frac{1}{2}) \left(\frac{\lambda}{G_{S,1Q}} + 1 - \lambda \right) & 0 & 0 \\ 0 & n_T + \frac{1}{2} & 0 \\ 0 & 0 & n_A + \frac{1}{2} \end{bmatrix}, \quad (5.41)$$

where both the partially attenuated squeezed vacuum term ($\propto \lambda/G_{S,1Q}$), and the unsqueezed contribution from the loss mode ($\propto 1 - \lambda$) appear clearly. Note that in the absence of loss,

$$\Sigma_{i,\mathbf{X}} = \mathbf{S}_{\mathbf{X}}^* \Sigma_{\text{in},\mathbf{X}} \mathbf{S}_{\mathbf{X}}^{\text{T}}, \quad (5.42)$$

in agreement with Eq. (5.26).

We can now calculate $\Sigma_{\text{out},\mathbf{X}}$, the output spectral density matrix along \vec{x} . Given Eq. (5.21) for the cavity's susceptibility in the quadrature basis we may write

$$\Sigma_{o,\mathbf{X}} = \chi^*(\omega) \Sigma_{i,\mathbf{X}} \chi(\omega)^{\text{T}}. \quad (5.43)$$

Then losses between cavity and AMP are accounted for in the same manner as before, and finally the AMP performs another OMS operation, amplifying the SQ squeezed quadrature. We thus obtain

$$\begin{aligned} \Sigma_{\text{out},\mathbf{X}} &= \mathbf{A}_{\mathbf{X}}^* [\mathbf{L}_{\mathbf{X}}^* \Sigma_{o,\mathbf{X}} \mathbf{L}_{\mathbf{X}}^{\text{T}} + \mathbf{N}_{\mathbf{X}}] \mathbf{A}_{\mathbf{X}}^{\text{T}} \\ &= \mathbf{A}_{\mathbf{X}}^* [\mathbf{L}_{\mathbf{X}}^* \chi^*(\omega) \Sigma_{i,\mathbf{X}} \chi(\omega)^{\text{T}} \mathbf{L}_{\mathbf{X}}^{\text{T}} + \mathbf{N}_{\mathbf{X}}] \mathbf{A}_{\mathbf{X}}^{\text{T}} \end{aligned} \quad (5.44)$$

as the analog of Eq. (5.26) in the presence of loss.

The first entry in the matrix $\Sigma_{\text{out},\mathbf{X}}$ is the output spectral density at the measurement port along \hat{X} :

$$\begin{aligned} \Sigma_{\text{out},\mathbf{X},m} &= \left(n_T + \frac{1}{2} \right) G_{A,1Q} (1 - \lambda) + \frac{G_{A,1Q} \lambda}{B(\omega)} \left[\left(n_A + \frac{1}{2} \right) \kappa_a \kappa_m \right. \\ &\quad \left. + \left(n_T + \frac{1}{2} \right) \left(\kappa_l \kappa_m + \left(1 - \lambda + \frac{\lambda}{G_{S,1Q}} \right) \beta(\omega) \right) \right], \end{aligned} \quad (5.45)$$

where $B(\omega)$ and $\beta(\omega)$ are defined as in Eq. (5.30), and the same approximations have been made.

From $\Sigma_{\text{out},\mathbf{X},m}$ we can extract the signal visibility $\alpha(\omega)$ in the presence of loss:

$$\alpha(\omega) = \frac{\lambda n_A}{\left(n_T + \frac{1}{2} \right)} \times \frac{\kappa_a \kappa_m}{B(\omega) (1 - \lambda) + \lambda \left[\kappa_l \kappa_m + \left(1 - \lambda + \frac{\lambda}{G_{S,1Q}} \right) \beta(\omega) \right]}. \quad (5.46)$$

When $\lambda = 1$, Eq. (5.46) reduces to Eq. (5.5).

Finally, the scan rate enhancement E_t in the presence of loss, presented in Fig. 5.2b, is $\int_{-\infty}^{\infty} \alpha^2(\omega) d\omega$, normalized by the same integral with $G_{S,1Q} = 1$ and $\kappa_m = 2\kappa_l$.

5.1.3 Equivalence between single- and double-quadrature amplification

Single-mode squeezing can only enhance sensitivity to displacements along a single quadrature of the cavity field. In a situation such as an axion search, the signal distributes its excess power equally between the two quadratures. Thus switching to single-quadrature measurement from double-quadrature measurement, currently the operational mode of choice of axion haloscopes [105, 106], seems detrimental. In this section, we will see that in the absence of squeezing, neither measurement scheme has an advantage over the other.⁹⁷

If we neglect amplifier added noise, the signal visibility α is independent of whether we measure one quadrature or both. Specifically, an axion signal characterized by its spectral density S_a at the amplifier input divides itself equally as $S_a/2$ between the two quadratures. Similarly, the vacuum noise spectral density $\hbar\omega_a/2$ divides its power equally as $\hbar\omega_a/4$ between the quadratures.

However, we must account for the quantum limits on two-quadrature measurements. Any linear amplifier that measures both quadratures adds at least a second half-quantum of input-referred noise, evenly distributed between the two quadratures [128]. An ideal double-quadrature measurement thus yields

$$\alpha_{2Q} = \frac{S_a}{\hbar\omega_a} \quad (5.47)$$

in each quadrature. By comparison, there is no quantum limit on single-quadrature amplification, so an ideal single-quadrature measurement yields

$$\alpha_{1Q} = \frac{2S_a}{\hbar\omega_a} = 2\alpha_{2Q} \quad (5.48)$$

in the amplified quadrature.

To make a fair comparison between the single-quadrature and double-quadrature cases, we must consider the improvement in the SNR (defined in Sec. 5.1.1). Because all pertinent differences between the two cases enter when considering a single tuning step, we neglect tuning in the following

⁹⁷ This will be true insofar as there is negligible added noise at all frequencies, and no bandwidth limitations in either case. Johnson-Nyquist noise from an identical source is assumed in order to make an apples-to-apple comparison.

discussion. The SNR $\bar{\alpha}$ is given in terms of the spectral density ratio α by

$$\bar{\alpha} = \sqrt{\frac{N}{2}} \frac{\Delta_a}{\Delta} \alpha, \quad (5.49)$$

where Δ_a/Δ is the ratio of signal to noise bandwidths, and N is the number of independent measurements of the voltage contained in the noise bandwidth Δ . In considering the appropriate values of N and Δ for the two cases of interest, we will find two independent effects, each reducing $\bar{\alpha}_{1Q}$ by a factor of $\sqrt{2}$ relative to $\bar{\alpha}_{2Q}$. Together, these effects cancel the apparent benefit stemming from the absence of quantum noise limits in the single-quadrature case.

First, the Nyquist theorem guarantees that there are

$$N_{2Q} = 2\tau\Delta_{2Q} \quad (5.50)$$

independent measurements of the noise voltage in a double-quadrature measurement of duration τ and bandwidth Δ_{2Q} , where the factor of 2 counts the two independent quadrature amplitudes for each resolved Fourier component. Thus there are

$$N_{1Q} = \tau\Delta_{1Q} \quad (5.51)$$

measurements of the noise voltage in a single-quadrature measurement of bandwidth Δ_{1Q} .

Second, noiseless single-quadrature measurement with a parametric amplifier creates an irreversible ambiguity between the output signal and idler frequencies, equally spaced about the amplifier band-center. This ambiguity necessitates mapping amplifier outputs at a given detuning from band-center to the input signal and idler frequencies (Sec. 5.7). The consequent addition of spectral densities, half of which are guaranteed not to have an axion-induced excess power, effectively increases the noise bandwidth in a single-quadrature measurement by a factor of 2 relative to the signal bandwidth: $\Delta_{1Q} = 2\Delta_a$. In the standard double-quadrature measurement scheme, the signal and noise bandwidths are equal: $\Delta_{2Q} = \Delta_a$.

Putting this all together, the two measurement schemes are seen to be equivalent:

$$\begin{aligned}
\bar{\alpha}_{1Q} &= \sqrt{\frac{\tau\Delta_{1Q}}{2}} \frac{\Delta_a}{\Delta_{1Q}} \alpha_{1Q} \\
&= \sqrt{\frac{2\tau\Delta_{2Q}}{2}} \frac{\Delta_a}{2\Delta_{2Q}} 2\alpha_{2Q} \\
&= \bar{\alpha}_{2Q}.
\end{aligned} \tag{5.52}$$

The first line agrees with Eq. (5.33) for the single-quadrature SNR. The advantage of squeezed over unsqueezed single-quadrature haloscope operation can thus be taken as the true advantage that squeezing confers over current haloscopes that amplify both quadratures.

5.2 Model for the axion field

Figure 5.1 models the axion field as a fictitious generator that drives the cavity through a weakly coupled port. In this section, we relate the fictitious port coupling κ_a and the power spectral density n_A at the generator output to physical parameters normally found in the haloscope literature, and show that for representative values, the axion field acts as a classical force.

The measurable axion-sourced power is obtained from the first term in the output spectral density, Eq. (5.30). Referring to the cavity output and multiplying by the axion linewidth Δ_a , the on-resonance ($\omega_a = \omega_c$) signal power is

$$P_{\text{sig}} = 4\hbar\omega_a n_A \Delta_a \frac{\kappa_a \kappa_m}{(\kappa_m + \kappa_l)^2}. \tag{5.53}$$

The two model parameters κ_a and n_A can readily be related to physical parameters by comparing Eq. (5.53), to e.g. Eq. (1) of Ref. [90] evaluated on resonance. In our notation this expression takes the form

$$P_{\text{sig}} = \left(g_{a\gamma\gamma}^2 \frac{\hbar c^3 \rho_a}{\mu_0 \omega_a^2} \right) \left(B_0^2 V C_{mnl} \frac{\omega_c^2 \kappa_m}{(\kappa_m + \kappa_l)^2} \right), \tag{5.54}$$

where μ_0 is the vacuum permeability and other parameters are as defined in Eqs. (2.8) and (2.16).

Equating Eqs. (5.53) and (5.54), we obtain

$$n_A \kappa_a = \frac{g_{a\gamma\gamma}^2 \rho_a c^3}{4\omega_a \mu_0 \Delta_a} B_0^2 C_{mnl} V. \tag{5.55}$$

To derive a second expression relating κ_a to n_A , we observe that the fictitious generator in our haloscope model may equivalently be represented as a second harmonic oscillator mode with very high occupancy but very weak coupling to the haloscope cavity. This equivalence is illustrated in Fig. 5.3. The axion field in any laboratory-scale volume constitutes such an oscillator mode, with resonant frequency ω_a . Specifically, we model the oscillations of the axion field as a fictitious cavity occupying the same volume as the real haloscope cavity. The quanta of this fictitious cavity are axion particles, so its total occupancy is

$$N_A = \frac{V\rho_a}{\hbar\omega_a}; \quad (5.56)$$

it is coupled to the haloscope cavity with interaction Hamiltonian

$$\hat{H}_{\text{int}} = \hbar g(\hat{A} + \hat{A}^\dagger)(\hat{B} + \hat{B}^\dagger), \quad (5.57)$$

where g is the interaction strength and \hat{A} (\hat{B}) is the annihilation operator of the haloscope (fictitious) cavity. The Hamiltonian of the closed system is

$$\hat{H} = \hat{H}_0 + \hat{H}_{\text{int}}, \quad (5.58)$$

where

$$\hat{H}_0 = \hbar\omega_c \left(\hat{A}^\dagger \hat{A} + \frac{1}{2} \right) + \hbar\omega_a \left(\hat{B}^\dagger \hat{B} + \frac{1}{2} \right). \quad (5.59)$$

Coupling the haloscope cavity to measurement and loss ports at rates κ_m and κ_l , respectively, we write down the Heisenberg-Langevin equations of motion for the open system:

$$\frac{d\hat{A}}{dt} = -i\omega_c \hat{A}(t) - ig[\hat{B}(t) + \hat{B}^\dagger(t)] - \frac{\kappa_m + \kappa_l}{2} \hat{A} + \sqrt{\kappa_m} \hat{a}_{\text{in},m}(t) + \sqrt{\kappa_l} \hat{a}_{\text{in},l}(t) \quad (5.60)$$

$$\frac{d\hat{B}}{dt} = -i\omega_a \hat{B}(t) - ig[\hat{A}(t) + \hat{A}^\dagger(t)], \quad (5.61)$$

where we describe the input bath associated with the measurement (loss) port with the annihilation operator $\hat{a}_{\text{in},m}$ ($\hat{a}_{\text{in},l}$).

We restrict ourselves to the classical limit of the system, with operators demoted to complex amplitudes, and the case where the resonances coincide, $\omega_a = \omega_c$, with no power entering via the




Figure 5.3: (a) Model for the haloscope cavity (with mode \hat{A}) corresponding to Fig. 5.1, wherein an axion signal is modeled as a microwave generator sourcing mean photon number n_A and coupled to the haloscope cavity at power decay rate κ_a . (b) Equivalent representation of the axion signal as a second, fictitious cavity (mode \hat{B}) filled with N_A axions coupled to the haloscope cavity with field exchange rate g . The other couplings of the haloscope cavity remain unmodified.

quantity	value
ρ_a	0.45 GeV/cm^3
B_0	9 T
$g_{a\gamma\gamma}$	$-7.7 \times 10^{-24} \text{ eV}^{-1}$
Δ_a	5 kHz
$\omega_a/2\pi$	5 GHz
V	1.5 L
C_{mnl}	0.5

Table 5.1: Representative physical values used to estimate axion port parameters.

measurement and loss ports,

$$a_{\text{in},m} = a_{\text{in},l} = 0. \quad (5.62)$$

Transforming into the rotating frame of the haloscope cavity,

$$A(t) \rightarrow A(t)e^{-i\omega_c t} \quad (5.63)$$

$$B(t) \rightarrow B(t)e^{-i\omega_c t} \quad (5.64)$$

and making a rotating wave approximation, Eqs. (5.60) and (5.61) reduce to

$$\frac{dA}{dt} = -igB(t) - \frac{\kappa_m + \kappa_l}{2}A(t) \quad (5.65)$$

$$\frac{dB}{dt} = -igA(t), \quad (5.66)$$

where g plays the role of the field exchange rate corresponding to the power decay rate out of the axion cavity,

$$g = \frac{\kappa_a}{2}. \quad (5.67)$$

These equations of motion describe an exchange of energy between the two cavities and a decay of that energy out from the haloscope cavity via the measurement and loss ports. We are interested in the steady-state ($dA/dt = 0$) output field

$$A_{\text{out},m} = -\sqrt{\kappa_m}A \quad (5.68)$$

when the occupancy of the axion cavity is

$$|B|^2 = N_A. \quad (5.69)$$

We find a steady-state occupancy of the haloscope cavity

$$|A|^2 = \left(\frac{\kappa_a}{\kappa_m + \kappa_l} \right)^2 N_A, \quad (5.70)$$

implying an output signal power of

$$P_{\text{sig}} = \hbar\omega_a |A_{\text{out},m}|^2 = \frac{\hbar\omega_a N_A \kappa_a^2 \kappa_m}{(\kappa_m + \kappa_l)^2}. \quad (5.71)$$

Eqs. (5.53) and (5.71) for the output power must agree, implying

$$\frac{n_A}{\kappa_a} = \frac{N_A}{4\Delta_a} = \frac{V\rho_a}{4\hbar\omega_a\Delta_a}. \quad (5.72)$$

In terms of physical parameters, Eqs. (5.55) and (5.72) yield

$$n_A = \frac{|g_{a\gamma\gamma}|\rho_a B_0 V}{4\omega_a\Delta_a} \sqrt{\frac{C_{mnl}c^3}{\hbar\mu_0}} \quad (5.73)$$

$$\kappa_a = |g_{a\gamma\gamma}|B_0 \sqrt{\frac{C_{mnl}\hbar c^3}{\mu_0}}. \quad (5.74)$$

Using representative values for the HAYSTAC experiment [90] in the presence of a 5 GHz KSVZ axion shown in Table 5.1, we obtain the values for our model parameters shown in Table 5.2. We see that the fictitious generator is well into the classical regime,

$$n_A \gg \frac{1}{2}, \quad (5.75)$$

while its extremely feeble coupling κ_a nonetheless makes its presence a challenge to detect.

parameter	value
N_A	3.3×10^{16}
$\kappa_a/2\pi$	2.3 μHz
n_A	2.4×10^7

Table 5.2: Model parameter values calculated using the physical values from Table 5.1.

5.3 Enhanced signal visibility with squeezing

In practice, the scan rate of Eq. (5.37) can be degraded not only by transmission losses, but also by distortions in the squeezed state (Ch. 4) and added noise in the amplifier chain. In order to investigate how accurately Eq. (5.37) predicts the performance of a squeezing-enhanced haloscope, we work with an apparatus that mimics its microwave network, but without some of the cumbersome features. Specifically, in an axion haloscope the mechanically tunable cavity must reside in a large static magnetic field to enhance κ_a and n_A . In our apparatus, the cavity has a fixed frequency $\omega_c = 2\pi \times 7.146$ GHz and there is no applied field. Consequently the cavity, constructed from superconducting aluminum and shown in Fig. 5.4, has much lower intrinsic loss than a copper haloscope cavity. Additional loss is added through an explicit port that extracts energy from the cavity at a rate $\kappa_l = 2\pi \times 100$ kHz, a typical value for a haloscope cavity. An axion-like signal is coupled into the cavity through a microwave generator connected to a second port weakly coupled at rate $\kappa_a = 2\pi \times 100$ Hz. Finally, a third port couples the cavity mode to an SSR with a rate chosen to be close to the optimum value for the case with ($\kappa_m = 10\kappa_l$) or without ($\kappa_m = 1.5\kappa_l$) squeezing, creating a physical realization of the model in Fig 5.1.

For the largest increase in scan rate, the SSR should be attached to the cavity measurement port with as little transmission loss as possible. To investigate the transmission loss independent of the cavity loss, we use the fact that the JPAs are narrowband (~ 5 MHz), tunable amplifiers and detune both the SQ and AMP far off cavity resonance (~ 10 MHz) so that the squeezed state is promptly reflected from the cavity. Figure 5.5 illustrates the ability to efficiently generate, transport, and amplify a squeezed state in this off-resonance configuration. It shows histograms of the measured voltage in the AMP's amplified quadrature $\hat{X}_{\text{out},m}$ as a function of θ , the phase




Figure 5.4: (a) Photo of the aluminum cavity (closed up) used for the experiments in Secs. 5.3 and 5.4. The three SMA ports are for measurement, loss, and signal injection. A penny is shown for scale. (b) Same cavity, opened up. The 7.146 GHz resonant mode lives in the slot visible in both halves of the cavity.

between the amplified quadrature of the SQ and the amplified quadrature of the AMP. When $\theta = \pi/2$ or $3\pi/2$, one quadrature is first squeezed and then amplified. At these points, comparing the output noise variances σ_{on}^2 and σ_{off}^2 measured with SQ on and off (not pumped), respectively, the squeezing $S = -4.5 \pm 0.1$ dB using the definition of Eq. (4.4). As in the last chapter (Footnote 90), this is not an inferred squeezing; rather, it is a direct measurement of an overall 4.5 dB reduction in the noise floor of $\hat{X}_{\text{out},m}$ over the whole bandwidth of the quadrature measurement. This amount of squeezing (consistent with the estimate of $\eta = 0.69 \pm 0.01$, Sec. 5.5.3) required particular care in reducing the transmission losses between the two JPAs, which was facilitated by using the flux-pumped JPAs of Sec. 3.7.

This large amount of delivered squeezing implies that the SSR should improve our ability to resolve a weakly coupled signal detuned from cavity resonance. To demonstrate this improvement, the SQ and AMP bands are centered on cavity resonance with the cavity's measurement port overcoupled ($\kappa_m = 10\kappa_l$), and the phase θ set to $\pi/2$ or $3\pi/2$ (Fig. 5.5). To perform the single-quadrature measurement, the output $\hat{X}_{\text{out},m}$ is mixed down with a local oscillator (LO) at the cavity resonant frequency its power spectral density is computed. As for any such measurement, the frequency component at ω in the down-mixed output is a linear combination of two frequency components at the mixer's high frequency input, $\pm\omega$ detuned from the LO frequency. Figure 5.6a shows the spectral density when the tone is 1 MHz detuned from the cavity's resonance. Comparing




Figure 5.5: (a) Histograms of voltage fluctuations V_X measured along $\hat{X}_{\text{out},m}$, as a function of the SQ-AMP relative phase θ . (b) Corresponding vacuum squeezing $S = \sigma_{\text{on}}^2/\sigma_{\text{off}}^2$. Data taken off of cavity resonance. Figure reproduced from Ref. [208].

two situations, SQ on and SQ off, the signal visibility improves by roughly 4 dB in the presence of squeezing, with 0.5 dB degradation from cavity loss.

To estimate the associated increase in scan rate from the SSR, the tone is stepped across the cavity's resonance and the visibility $\alpha(\omega)$ measured at each detuning ω . In order to compare the optimal squeezed and unsqueezed cases, Fig. 5.6b shows measured values of $\alpha(\omega)$ for two cases: for $\kappa_m = 10\kappa_l$ with squeezing, and for $\kappa_m = 1.5\kappa_l$ without squeezing (see Sec. 5.6 for two other complementary cases). Without squeezing, $\alpha(0)$ is greater but the bandwidth is poor. With squeezing, α remains relatively high as ω increases. Extracting $R \propto \int_{-\infty}^{\infty} \alpha(\omega)^2 d\omega$ for the two cases, the estimated scan rate enhancement is $E_e = 2.05 \pm 0.04$. Independently measured values for η , κ_m , and $G_{S,1Q}$, yield expected values for $\alpha(\omega)$, shown as solid lines, in excellent agreement with the measured values in Fig. 5.6b. Finally, from the expected $\alpha(\omega)$, $E_t = 2.11 \pm 0.07$, also in quantitative agreement with the data-based estimate E_e .




Figure 5.6: (a) Improvement in microwave tone visibility due to squeezing. Power spectra normalized to the unsqueezed vacuum power, with (red) and without (purple) squeezing are shown in for a tone 1 MHz detuned from the resonant frequency of the overcoupled cavity ($\kappa_m = 10\kappa_l$). The x -axis has been shifted between the two situations for visual clarity. (b) Visibility $\alpha(\omega)$ measured as a function of the tone’s detuning from cavity resonance ($\omega = 0$) for two cases: no squeezing while near critical coupling ($\kappa_m = 1.5\kappa_l$, blue circles), and squeezing while strongly overcoupled ($\kappa_m = 10\kappa_l$, red circles). In both cases $\alpha(\omega)$ is normalized to the expected maximum value $\alpha_{\max} = \alpha(0)$ evaluated at $\kappa_m = \kappa_l$. The theoretical expectation for each case is superimposed as a black curve, calculated using $\kappa_l = 2\pi \times 100$ kHz, $\eta = 0.69$ and $G_{S,1Q} = 13$ dB. Figure reproduced from Ref. [208].

5.4 Mock-haloscope experiment: accelerated search for a weak classical tone

In a real axion search, the aim is to detect a signal many orders of magnitude smaller than that in Fig. 5.6. Inferring the presence or absence of an axion at each frequency requires combining the measured powers from many adjacent cavity tunings [78, 125]. Furthermore, over a long integration time the benefit inferred from Fig. 5.6 is vulnerable to practical non-idealities. Drifts of either JPA’s gain, drifts of the SQ-AMP relative phase θ , non-Gaussian noise processes, and interfering radio frequency (RF) or intermediate frequency (IF) tones are of particular concern. In this section, we will see that the SSR indeed matches the performance presented in the previous section when searching for a feeble tone over a wide frequency range.⁹⁸

In order to fully validate the enhancement gained from the SSR, a fake axion, or “faxion,”

⁹⁸ This experiment was informally given the name “the Dummy Experiment for Axionic Dark Matter Active Upwards of 5 GHz,” or DEADMAU5.

tone is sent through the cavity weakly coupled port in order to be detected. The tone is synthesized from a randomly modulated microwave tone, whose power is adjusted such that the faxion spectral density emerging from the cavity is roughly 1% of vacuum, and whose width is broadened to roughly $\Delta_a \approx 9$ kHz, comparable to expectations for a realistic axion at frequency $\omega_a \approx 2\pi \times 7$ GHz (Sec. 2.4.3). Stepping the faxion tone frequency “backwards” past a stationary cavity simulates a realistic axion search without the hardware demands imposed by a tunable cavity. The faxion’s initial frequency is chosen randomly within a 2 MHz window around the cavity resonance, and is then tuned in discrete -10 kHz steps over a 4 MHz window. At each tuning step, an output power spectral density is recorded as in Sec. 5.3. These spectra are mixed down and referred to the mixer’s input by symmetrizing about $\omega = 0$. Spectra are artificially shifted in steps of +10 kHz to simulate cavity tuning. The spectra are then rescaled by $\alpha(\omega)$ such that frequency bins with higher sensitivity to the faxion are weighted more. These rescaled, shifted spectra are then added into a grand spectrum (Sec. 5.7 and Ref. [78]). With this procedure, the faxion’s contributions in each individual spectrum add at its initial frequency as if the cavity had been tuned, creating a clear excess of power in the grand spectrum. Figure 5.7a presents some symmetrized spectra, obtained while squeezing with an overcoupled cavity ($\kappa_m = 10\kappa_l$), in which the faxion excess power is too small to rise above vacuum fluctuations. The spectra were obtained at different fictitious cavity tunings, normalized to their measured standard deviations, and vertically offset from one another for visual clarity. In the resulting grand spectrum (Fig. 5.7b), a prominent faxion peak emerges 6σ above the noise.

Extracting a scan rate enhancement from such a realistic signal search requires two distributions of faxion powers: one when squeezing and one when not squeezing, with near-optimal κ_m for each case. Each distribution is obtained by repeating the faxion injection and detection protocol 200 times. Over the course of the measurement, which takes roughly 9 hours per configuration, the relative phase θ between SQ and AMP quadratures is stabilized at $\pi/2$ via a feedback loop (Sec. 5.5.2). Figure 5.7c displays the two faxion power distributions, as well as the noise power




Figure 5.7: (a) Processed spectra. The acquired power spectra are symmetrized, rescaled and then shifted in frequency to align all the frequency bins containing the faxon with each other (red dots), thereby effectively tuning the cavity. The power excess of the processed spectra is plotted in units of their standard deviation σ^γ . (b) Grand spectrum constructed from all processed spectra. The grand spectrum excess power is plotted in units of its standard deviation σ^η . A large power excess is observed at the frequency of the faxon. (c) Distributions of faxon powers for many separate experiments each culminating in a grand spectrum. When squeezing with $\kappa_m = 10\kappa_l$, the powers (red points) are drawn from a Gaussian distribution $\mathcal{N}(\mu_s, (\sigma^\eta)^2)$ (solid line), with $\mu_s = 6.05 \pm 0.07$. When not squeezing with $\kappa_m = 1.5\kappa_l$, the powers (blue points) are drawn from $\mathcal{N}(\mu_u, (\sigma^\eta)^2)$ (solid line), with $\mu_u = 4.15 \pm 0.07$. The powers (black points) drawn from the faxon-less noise distribution $\mathcal{N}(0, (\sigma^\eta)^2)$ (solid line) are also represented for comparison. Figure adapted from Ref. [208].

distribution. The no-faxon distribution of mean zero and variance $(\sigma^\eta)^2$,⁹⁹ Eq. (6.12), is guaranteed by the central limit theorem to be Gaussian distributed, $\mathcal{N}(0, (\sigma^\eta)^2)$. The faxon adds a small mean excess of power μ insufficient to enlarge the variance. The faxon distribution is therefore $\mathcal{N}(\mu, (\sigma^\eta)^2)$. The signal and noise distributions separate as the total measurement time squared,

⁹⁹ For consistency with future chapters, the superscript η borrow from the notation of Sec. 6.1.

and thus the speed-up due to squeezing is equal to

$$E_m = \left(\frac{\mu_s}{\mu_u} \right)^2, \quad (5.76)$$

where μ_s (μ_u) is the mean of the faxion power distribution obtained when squeezing (not squeezing). The two distributions of faxion powers are characterized by $\mu_s = 6.05 \pm 0.07$ and $\mu_u = 4.15 \pm 0.07$, leading to a measured scan rate enhancement of $E_m = 2.12 \pm 0.08$, in quantitative agreement with the estimate E_e obtained from visibility measurements in Sec. 5.3.

This measurement provides the strongest proof out of all the arguments presented in this chapter that quantum squeezing not only can in principle be used to accelerate dark matter searches, but that it can do so subject to realistic transmission losses in a haloscope-like system. The following sections describe some of the auxiliary measurements and protocols necessary to perform the experiments in this chapter.

5.5 Experimental details and calibration

This section presents the experimental setup used for the experiments in Secs. 5.3 and 5.4, describes the phase-locking protocol used to ensure $\theta = \pi/2$ or $3\pi/2$ between SQ and AMP pump tones, and discussed the calibration measurement used to infer the intra-SSR transmission losses.

5.5.1 Experimental setup

The full experimental setup is represented in Fig. 5.8. The SQ-cavity-AMP ensemble is attached to the bottom plate of a dilution refrigerator. For each JPA, the pump tone is routed through a 10 dB directional coupler connected to the pump port. The placement and configuration of these directional couplers ensures that the large pump tone required for flux pumping primarily heats up a 50Ω termination whose Johnson-Nyquist noise [206, 207] propagates back up the pump line, away from the JPA. A coil around each chip, connected to a DC current source at room temperature, generates a DC magnetic field, and the chip-coil ensemble is magnetically shielded with aluminum and cryoperm. Each JPA is connected to a circulator through superconducting

Nb/Ti cables in order to minimize transmission losses.

An ensemble of four circulators route the squeezed state and provide microwave isolation between the SSR elements. Two circulators between the SQ and the cavity protect the SQ from power reflected back from the cavity. These circulators provide sufficient isolation when the SQ is operated with 13 dB of signal gain. Similarly, three circulators separate the AMP and the cavity, as the AMP is operated with higher signal gain, around 25 dB. Experimentally, it was observed that with only two circulators, undesirable feedback between the cavity and AMP perturbs the AMP's gain by effectively changing its pump's power.

At the chain's input, either vacuum noise or a probe tone can be injected via a 20 dB directional coupler. This probe tone is useful when characterizing the JPA gain profiles or when characterizing the cavity. Finally, at the output, a double-junction isolator protects the AMP from signals reflected from the next amplifier, a high-electron-mobility transistor (HEMT) amplifier at 4 K.

At room temperature, a single microwave generator (Keysight E8257D) drives both JPAs and also serves as LO for the in-phase/quadrature (IQ) mixer in the readout line. On the path leading to the SQ pump's input, a voltage-controlled variable attenuator and phase shifter provide control of both the SQ pump amplitude and phase. On the path leading to the IQ mixer, a frequency divider converts ω_p to ω_0 , the bare resonant frequency of both JPAs and the cavity. This divider (Pasternack PE88D2000) can input a wide range of powers (from -20 to 5 dBm) while always outputting the same power (roughly -4 dBm). Thus, the AMP gain can be tuned freely with the microwave generator's output power. A second generator injects a tone through the weakly coupled port of the axion cavity. This tone can be shaped into a 9 kHz wide Lorentzian with an arbitrary waveform generator (AWG).

At the output, a 6 dB directional coupler routes a fraction of the power to a vector network analyzer (VNA), which monitors signals either from the SSR chain's input or from the cavity's weakly coupled port. The VNA is used to measure *in situ* the couplings κ_m and $\kappa_l + \kappa_a$. The other portion of the output power reaches the IQ mixer's RF port. After being mixed down, the in-phase

and quadrature signals are amplified and directed onto 1.9 MHz anti-aliasing low-pass filters, then finally digitized by an analog-to-digital converter (ADC).

5.5.2 Phase-locking of SQ and AMP pumps

Enhancing the scan rate with squeezing is only beneficial if the AMP amplifies the SQ squeezed quadrature. To this end, as suggested by Fig. 5.5, the phase θ between SQ and AMP pumps must be maintained at $\pi/2$ (or $3\pi/2$) to a high precision, given the sharp dependence of the degree of squeezing S on θ . When $S > 0$ dB, the output is noisier with the SQ on than off, and the SSR is therefore detrimental to the axion search. Furthermore S also depends on the SQ gain $G_{S,1Q}$, i.e. on the SQ pump power, and there is an optimal $G_{S,1Q}$ for which S reaches its minimum value S_{\min} . In fact, as the gain increases, S first improves as the squeezed state elongates in phase space, but then saturates due to distortion effects (Ch. 4). For this type of JPA, $S = S_{\min}$ at $G_{S,1Q} \approx 13$ dB is determined experimentally.

A feedback loop that uses the output variance σ_{on}^2 as its control parameter on the voltage-controlled variable attenuator and phase shifter is used in order to achieve $S = S_{\min}$ throughout a data run. When initializing a 9 hour acquisition, a two-dimensional sweep of the variable attenuation A_s and phase shift θ is used to estimate the global minimum of σ_{on}^2 . Then, periodically throughout the data run, a fast gradient descent-type algorithm corrects for small drifts of σ_{on}^2 . Note that this approach is robust to possible phase shifts due to changes in the variable attenuation and vice versa. Empirically, S can be maintained at $S_{\min} \pm 0.1$ dB, over the course of the entire experiment.

The AMP pump power remained stable around 25 dB without adjustment. However, when implementing the SSR in HAYSTAC (Ch. 8), it has to be adjusted because the frequency of the JPAs is also stepped in time. Note that in a practical haloscope run, sizable fluctuations of the net receiver gain in a SSR-integrated setup on timescales shorter than the raw spectrum acquisition time can be detrimental to axion detection. This is true even if all sources of added noise are overwhelmed. For our raw spectra acquisition times of 0.32 s, the receiver gain fluctuations are




Figure 5.8: Full experimental setup used for the mock-haloscope experiment. The SSR is shown together with its room temperature control electronics. The input (black) and output (red) coaxial cables are in bold. At the 20 mK stage, Nb/Ti superconducting coaxial cables are shown in green, and the three adjacent circulators form a triple-junction circulator fabricated by QuinStar (QCE-070100CM30). The thinner (orange) lines represent DC cables. Figure reproduced from Ref. [208].

negligibly small.

5.5.3 Intra-SSR loss measurement

The presence of four circulators and several microwave connectors, including adapters from SMP (used on the JPAs chips) to SMA standards, inevitably reduces the transmission efficiency η between SQ and AMP. A triple-junction circulator¹⁰⁰ and superconducting SMA cables between the two JPAs and the circulators are used to minimize η . However, η still provides the primary limitation on the efficacy of the SSR in accelerating axionic dark-matter searches.

Measuring the output power spectral density $P_{S/A}^O$ of vacuum fluctuations in a single quadrature amplified by the SQ with the AMP off then repeating this measurement with the role of the JPAs interchanged yields an estimate of η . Performing the exact measurement described in Sec. 4.4.2 gives $\eta = G_{\text{SQ}}^O/G_{\text{AMP}}^O = 0.69 \pm 0.01$. Figure 5.9 shows $4P_{S/A}^O/(\hbar\omega W)$, when varying either the SQ or the AMP gain. We extract $\eta = G_{\text{SQ}}^O/G_{\text{AMP}}^O = 0.69 \pm 0.01$.

The efficiency η observed in the mock-haloscope setup should not be significantly degraded by the large magnetic fields required in a real haloscope experiment such as HAYSTAC or ADMX, as is borne out in Ch. 8. This is because the primary constraint imposed by the large field is an increased spatial separation between the axion cavity and the circulators and amplifiers that process the signal. This ~ 1 m distance is bridged with superconducting coaxial cables, whose attenuation [232] is far subdominant to other sources of transmission loss, such as the microwave circulators in the setup. At sufficiently high fields, the cables would eventually experience significantly increased transmission losses. However, empirically, at the fields we operate in, the performance remains unaffected [90].

From η we can estimate the squeezing S that we should obtain far off cavity resonance (~ 10 MHz detuned). Considering the ideal case where the reduction of the squeezed state variance is

¹⁰⁰ Perhaps the most obvious upgrade available to this setup would be an L-shaped, cryogenic, quadruple-junction circulator to replace the combination of single- and triple-junction devices in Fig. 5.8. This option was not pursued here or in Ch. 8 purely on grounds of financial cost.




Figure 5.9: Intra-SSR transmission loss calibration measurement. The normalized single-quadrature output power spectral density $P_{S/A}^O$ is plotted as a function of the SQ or AMP gain $G_{A/S,1Q}$. Two configurations are represented: when operating only the SQ (red circles) and when operating only the AMP (blue circles). A linear fit (solid lines) to these linear responses allows yields η . Figure adapted from Ref. [208].

equal to the SQ single-quadrature power gain $G_{S,1Q}$, we obtain:

$$S = \frac{\eta}{G_{S,1Q}} + 1 - \eta, \quad (5.77)$$

which, for $G_{S,1Q} = 13$ dB leads to $S = -4.6$ dB, in quantitative agreement with what was measure in practice.

5.6 Complementary measurements of the microwave tone's improvement in resolution

We saw in Fig. 5.6b the visibility $\alpha(\omega)$ of a microwave tone as a function of ω , the detuning between the tone's frequency and the cavity bare resonance. We compared $\alpha(\omega)$ between the two relevant cases: squeezing with a strongly overcoupled cavity versus not squeezing with a near-critically coupled cavity. Figure 5.10 presents these, along with measurements of $\alpha(\omega)$ for the two




Figure 5.10: Complementary measurements of signal visibility in the SS. Signal visibility $\alpha(\omega)$ is plotted as a function of the signal's detuning from cavity resonance. It is normalized by α_{\max} , obtained at zero detuning for $\kappa_m = \kappa_l$. Four cases have been measured, with the corresponding theoretical predictions represented as solid lines: overcoupling ($\kappa_m = 10\kappa_l$) with squeezing (red circles), near-critically coupling ($\kappa_m = 1.5\kappa_l$) without squeezing (blue circles), overcoupling without squeezing (purple triangles), and near-critically coupling with squeezing (green triangles). The theory lines have been calculated for $\kappa_l = 100$ kHz, $\eta = 0.69$, and $G_{S,1Q} = 13$ dB. Figure reproduced from Ref. [208].

complementary cases: strongly overcoupling without squeezing, and near-critically coupling with squeezing. In all four cases, there is excellent agreement with predictions from the theory, developed in Sec. 5.1.2 and shown in solid lines. The theory curves are not fits; they use parameter values for κ_l , η , and $G_{S,1Q}$ measured independently. Furthermore, the values are the same across all four cases.

Squeezing is beneficial, even when not overcoupling the cavity's measurement port, as it enhances α off cavity resonance while having no effect on resonance. Experimentally, the complementary estimate $E_e^{c1} = 1.72 \pm 0.03$ is obtained when squeezing and near-critically coupling, compared to the same situation without squeezing, in agreement with the theoretical prediction, $E_t^{c1} = 1.77 \pm 0.03$. Conversely, the scan rate is worse when strongly overcoupling without squeezing, compared to near-critically coupling with squeezing: from the data, $E_e^{c2} = 0.57 \pm 0.01$, not far from

the theoretical value of $E_t^{c2} = 0.52$.

5.7 Processing of the SSR spectra

In order to detect the small cavity displacement of Sec. 5.4, 401 “raw” spectra are acquired, processed, and combined into one “grand” spectrum. The injected faxion tone is sufficiently small relative to the size of the vacuum fluctuations that in any one raw spectrum it typically does not stand out. However, after processing the spectra, the faxion emerges often well above the level of vacuum noise, as in Fig. 5.7b. This section provides a brief overview of the steps in the processing of the spectra and provides the specific numerical values used throughout the processing in order to produce Fig. 5.7.¹⁰¹ The processing procedure is based closely on the work of Ref. [78], which is described quantitatively in Sec. 6.1. This section borrows the terminology used for the intermediate processing stages set forth in Ref. [78], and also omits many details and rationales which are there covered extensively.

Using the setup of Fig. 5.1, fluctuations emerging from the cavity, possibly in the presence of squeezed vacuum noise, are directed into a chain of amplifiers led by the AMP. The fluctuations are mixed down to DC, further amplified, low-pass filtered up to 1.9 MHz, and sampled with an Alazar ATS9462 digitizer at 6 MS/s. Each raw spectrum, distinguished as the least processed data saved to a hard disk, itself actually comprises the frequency-averaged power spectral densities of 32 “subspectra,” each acquired over 10 ms and fast Fourier transformed to provide a spectral resolution of $\Delta_b = 100$ Hz. Since 401 raw spectra are acquired for each of the 200 squeezed and 200 unsqueezed data runs, the live acquisition time of all the spectra totals just over 14 hours, not counting the dead time. In practice it took roughly 18 hours.

Once the raw spectra are in hand, they are symmetrized. That is, the measured output at each signal frequency is processed as if it came equally from the signal and idler inputs. Since the output of a JPA at a given detuning from the center of its amplification band is nearly identical to the output at the opposite detuning from band center, Eqs. (3.128) and (3.129), the best one can

¹⁰¹ Some of these values differ than those used to process the HAYSTAC data.

do is to infer the spectral density measured in the homodyne configuration as coming equally from both sides of the pump.

The symmetrized spectra are then averaged according to their real frequencies (i.e. not accounting for the fictitious stepping of the cavity used to simulate a haloscope search) in order to detect excess power in either the pertinent RF or IF band.¹⁰² This mean spectrum is then high-pass filtered by the profile of a Savitsky-Golay (SG) filter [233, 234] applied to the average spectrum. A SG filter is simply a computationally quick means of applying the d^{th} -degree polynomial generalization of a moving average within a window of $2W + 1$ bins. If the width W is much larger than the size of interesting spectral features, those features will be minimally attenuated when the spectrum is divided by the filter response, whereas features extending over wide bandwidths will be effectively removed. Following Ref. [78], $d = 10$, $W = 500$ for this first filtering.

The resulting real-frequency-averaged spectrum is used to detect and remove excesses of power which do not act like our faxion. Bins that exhibit power fluctuations more than 4 standard deviations above the mean at their real frequencies are discarded along with their close neighbors as contaminated for all spectra. In practice, over 99% of bins are preserved, with only extreme outliers eliminated. Note that since in the real frequency space the faxion tone is being stepped, its 400 tunings do not combine over any one narrow block of bins, and this bin rejection procedure will therefore not eliminate a faxion signal. Once contaminated bins are removed, the SG filter is recalculated, as it will be slightly different without the presence of outlier bins, and reapplied to each spectrum individually.

This baseline removal leaves a set of nearly-flat “normalized” spectra, each centered near 1. Residual structure left from fluctuations in the overall receiver gain or its profile is then removed by dividing out the SG profile ($d = 4$, $W = 500$) of each individual spectrum. The mean value of 1 is then subtracted from each spectrum to form the “processed” spectra, several of which are shown in Fig. 5.7a.¹⁰³

¹⁰² In a real haloscope search, since the LO being used for homodyne measurement would be stepped along with the cavity, non-axion-induced power excesses in the RF are trickier to reject than their IF counterparts [78].

¹⁰³ For visual clarity, only a fraction of the bins are kept in Fig. 5.7a and b.

Next we must rescale the spectra to account for the varying sensitivity to a faxion tone as a function of its detuning from cavity resonance in any given spectrum. To form these “rescaled” spectra, the processed spectra are divided by the relative visibility profile of the squeezed state receiver: Eq. (5.46) with an additional term accounting for the small contribution of the HEMT’s added noise referred to the input of the AMP, as the weak frequency dependence of AMP gain over the bandwidth of the cavity in the presence of spectrally flat HEMT added noise contributes a correspondingly weak frequency-dependence to the noise. The rescaled spectra have the important property that a given power excess from the weakly coupled port of the cavity produces a constant expectation value of rescaled spectrum power excess, regardless of detuning from the center of the cavity. As a consequence, the variance of the power distribution for each bin grows with detuning from cavity resonance.

Aligning bins along the fictitious frequencies for which the stepped faxion tone would appear stationary, it is then possible to construct the maximum likelihood (ML) estimate of the power excess in each bin. Doing so yields the “combined” spectrum, a single spectrum whose bins contain as few as one (at the extreme edges) and as many as hundreds of contributions from individual rescaled spectra. The combined spectra bins each have a resolution of Δ_b , far below the linewidth Δ_a of the faxion. The power within non-overlapping sets of $n_c = 10 \ll \Delta_a/\Delta_b$ bins are thus averaged, to form the “rebinned” spectrum, which has $n_c\Delta_b = 1$ kHz spectral resolution.

Finally, overlapping sets of $n_r = 41 \gg \Delta_a/n_c\Delta_b$ rebinned spectrum bins are combined, accounting for the independently experimentally determined lineshape of the faxion, in order to produce the ML estimate of the faxion-shaped power centered on each bin. The resulting spectrum of estimated powers is the grand spectrum, shown renormalized to have mean power excess 0 and standard deviation 1 in Fig. 5.7b. In the grand spectrum, the faxion typically stands out well above the surrounding noise.

The acquisition and processing of sets of 401 raw spectra are repeated 200 times apiece for the optimal squeezed and unsqueezed cases. The two data sets require slightly different processing, chiefly because the visibility profile of the unsqueezed case does not include the frequency-dependent

contribution from the squeezing. This results not from finite bandwidth effects of the SQ, but from the fact that squeezed noise is preferentially absorbed near cavity resonance by the cavity's loss port, wherefrom it is replaced with unsqueezed vacuum.

Each of the 400 total acquisitions provides one measured faxion power. The powers are histogrammed for the optimal squeezed and unsqueezed cases, along with the far larger number of grand spectrum powers of bins not containing a faxion, in Fig. 5.7c. While the absolute mean values of the squeezed and unsqueezed faxion distributions carry little meaning, as they scale with the somewhat arbitrary choice of faxion power, their squared ratio, which would be preserved for a faxion of any power, gives the scan rate enhancement obtained from squeezing, as discussed in Sec. 5.4.

5.8 Concluding remarks on the SSR for enhancing the scan rate of real haloscope searches

In this chapter, we have seen both in theory and in practice that partially replacing the vacuum noise in an axion haloscope-like experiment with squeezed vacuum circumvents the standard quantum limit on the noise of the measurement apparatus, enabling high sensitivity over a broader bandwidth and a more rapid search for a weak classical tone. Squeezing moreover promises to move haloscopes into a qualitatively distinct design parameter space: whereas a quantum limited haloscope's scan rate plateaus with improving microwave transmission efficiencies η , a sub-quantum limited haloscope will benefit almost arbitrarily as microwave losses are further reduced. Thus, haloscopes can reap larger benefits from increased efficiency (for example, a tenfold scan rate enhancement at $\eta = 0.91$) as low-loss quantum technologies such as on-chip circulators and directional amplifiers mature [149, 218, 222].

In the following two chapters, we will turn our attention to the secondary focus of this thesis: the statistics of axion searches. We will return in Ch. 8 to the technique of this chapter, and combine it along with that of Ch. 7 to the HAYSTAC experiment.

Chapter 6

Axion haloscope processing and analysis

For this chapter and the next, we will turn our attention from quantum to classical information processing. This chapter details the processing protocol¹⁰⁴ and statistical analysis¹⁰⁵ framework used by HAYSTAC and other axion haloscopes prior to the work of this thesis. In the following chapter, I will present a Bayesian analysis framework tailored to axion dark matter searches. We will see that this framework improves significantly upon the standard analysis framework presented in Sec. 6.2 of this chapter, and validates the standard processing protocol presented in Sec. 6.1 as already near-optimal. Both this chapter and the next present the work of Ref. [126].

6.1 Haloscope processing overview

This section provides a quantitative overview of the steps of the processing protocol described in much greater detail in Ref. [78].¹⁰⁶ In Sec. 7.2 of the next chapter, we will evaluate the protocol for how well it preserves information content pertinent to the existence of the axion. This section

¹⁰⁴ This processing protocol was partially previewed in Sec. 5.7 for the specifics of the mock-haloscope experiment. It did not include an accounting for correlations necessary to accurately predict the power delivered by an axion, and did not cast the steps in terms of their mathematical operations. The description in this section is intended to be more general, more complete, more quantitative, and directly applicable to a full haloscope experiment.

¹⁰⁵ This thesis draws a distinction between what is meant by “processing” and “analysis.” “Processing” refers to any manipulation of a measured quantity from the haloscope in order to obtain a final pair of numbers quantifying the measured power excess and sensitivity at each frequency. The “analysis” comes after the processing, and refers to steps that intake that pair of numbers in order to perform a statistical inference about the presence or absence of the axion at each frequency.

¹⁰⁶ The HAYSTAC processing protocol in Ref. [78] grew out of that developed by ADMX [125]. In practice, active haloscopes all use similar data processing protocols, and the overview given in this section extends beyond what is done by HAYSTAC, though for other experiments specifics such as the use of the Savitsky-Golay filter to remove backgrounds may vary. At the time of writing this thesis, HAYSTAC is the only haloscope to have recently published a full description of its processing and analysis. The ABRACADABRA experiment [111], similar in some respects to a haloscope, has recently published its processing and analysis practices [75] as well.

focuses on the steps that transform each individual Fourier bin on its way to becoming part of the final, or grand, spectrum. The processing of real haloscope data also involves steps for dealing with practical nonidealities that are out of the scope of this section. We do not address issues such as the identification and removal of persistent but nonaxionic spectral features, or the details of the digital high-pass filtering of the data (see Sec. 5.7 or Ref. [78]).

The data from an axion haloscope arrives as a pair (one for each quadrature)¹⁰⁷ of discretely sampled, identically distributed time streams of measured voltages. The time streams are broken up into separate pieces of duration $1/\Delta_b$ and discretely Fourier transformed into so-called subspectra,¹⁰⁸ with frequency bin spacing Δ_b roughly two orders of magnitude narrower than the axion linewidth Δ_a . The subspectra are acquired in batches of N_b at each of N_t distinct haloscope tuning steps. The Fourier frequencies in each subspectrum can be mapped faithfully back to the N_f radio frequency (RF) frequencies¹⁰⁹ of the fluctuations within the haloscope cavity. Ultimately, each RF frequency is associated with a single bin in each of a large number of distinct subspectra. Therefore, the complex subspectrum voltages are denoted $v_{Q,ijk}^\alpha$, where $Q = X$ or Y denotes the quadrature, $i = 1, \dots, N_f$ indexes the RF frequency, $j = 1, \dots, N_t$ the haloscope tuning step (chronologically ordered), and $k = 1, \dots, N_b$ the subspectrum (also chronologically ordered) at a given tuning step. The Greek superscript (here α) will advance through the processing steps detailed below, with letters farther down the alphabet being used to represent spectra obtained later at later stages of data handling.

At each tuning step, sensitivity data are acquired (or borrowed from a nearby tuning) as well. The sensitivity data, in a simplified picture, fully specify the expected noise and possible axion signal power at each frequency bin i . At each tuning step and for each frequency, the sensitivity parameters η_{ij}^γ , defined as the ratio of signal power to noise power for an axion with $g_\gamma = 1$ that delivers 100% of its power at the i^{th} RF frequency, can then be calculated.

¹⁰⁷ It is theoretically equally efficient to perform single-quadrature measurement (Sec. 5.1.3), in which case there is only one quadrature time stream, but none of the conclusions of this section are altered.

¹⁰⁸ Really, the subspectra come when we compute the power spectral density, but we will refer to the voltages before that step as the subspectrum voltages.

¹⁰⁹ The number of RF frequencies in earlier processing steps N_f is traditionally larger than N , the number of frequencies in the grand spectrum, as a later data processing step will average sets of spectrally adjacent bins.

Using the terminology of Ref. [78], the subspectrum data are processed ultimately into grand spectrum data through the following steps:

- (1) The real and imaginary parts of each subspectrum quadrature voltage are summed in quadrature so as to obtain the quadrature powers $x_{Q,ijk}^\alpha$ within each bin:¹¹⁰

$$x_{Q,ijk}^\alpha = (\text{Im}[v_{Q,ijk}^\alpha])^2 + (\text{Re}[v_{Q,ijk}^\alpha])^2. \quad (6.1)$$

- (2) The quadrature powers are summed to obtain the total power x_{ijk}^α in each subspectrum bin:

$$x_{ijk}^\alpha = x_{X,ijk}^\alpha + x_{Y,ijk}^\alpha. \quad (6.2)$$

- (3) All subspectrum powers at each tuning step and frequency are averaged to obtain the raw spectrum powers:

$$x_{ij}^\beta = \frac{1}{N_b} \sum_{k=1}^{N_b} x_{ijk}^\alpha. \quad (6.3)$$

- (4) The expected noise power $\langle x_{ij}^\beta \rangle$, obtained from the application of one or more Savitzky-Golay (SG) filter (or more generally a low-pass digital filter) at each frequency, is divided from each raw spectrum to obtain the processed spectra as:¹¹¹

$$x_{ij}^\gamma = \frac{x_{ij}^\beta}{\langle x_{ij}^\beta \rangle} - 1. \quad (6.4)$$

The processed spectra (and subsequent spectra) are power excesses: that is, departures of measured powers from the values expected absent an axion.

- (5) The processed spectra are divided by the independently obtained sensitivities to obtain the rescaled spectra:

$$x_{ij}^\delta = \frac{x_{ij}^\gamma}{\eta_{ij}^\gamma}. \quad (6.5)$$

¹¹⁰ Equation (6.1) is for each quadrature individually. In the case of a heterodyne measurement such as that performed in Phase I of HAYSTAC [90, 106], which uses an in-phase/quadrature (IQ) mixer to move tones to the intermediate frequency (IF) band, image rejection is first performed in software in order to separate components coming from within the analysis band from those coming from outside of it. In practice, HAYSTAC Phase I achieved better than 20 dB rejection of the unwanted sideband [42].

¹¹¹ Dividing out a low-pass filter profile realizes a high-pass filter, which is what we ultimately want in order to preserve the (narrow) signature of the axion.

The rescaled spectra have the property that an axion with $g_\gamma = 1$ unrealistically depositing all of its power entirely within any one bin shifts the mean power excess for that bin up from 0 to 1. Less axion-sensitive bins are therefore drawn from distributions with larger variances $\text{var}(x_{ij}^\delta) = (\sigma_{ij}^\delta)^2$.

- (6) In each RF frequency bin i , the combined spectrum power excess x_i^ϵ is taken as the maximum likelihood (ML)-weighted sum of all rescaled spectra power excesses:

$$\frac{x_i^\epsilon}{(\sigma_i^\epsilon)^2} = \sum_{j=1}^{N_t} \frac{x_{ij}^\delta}{(\sigma_{ij}^\delta)^2}. \quad (6.6)$$

The variance of the i^{th} combined spectrum bin is obtained from the variance of all contributing rescaled spectrum bins as

$$(\sigma_i^\epsilon)^2 = \left(\sum_{j=1}^{N_t} \frac{1}{(\sigma_{ij}^\delta)^2} \right)^{-1}, \quad (6.7)$$

where $(\sigma_{ij}^\delta)^2 = \infty$ at tuning steps that do not contribute to a given frequency bin i .

- (7) Sets of $n_c \ll \Delta_a/\Delta_b$ adjacent combined spectrum power excesses are scaled by $n_c n_r$, defined below, and the rebinned spectrum power excesses x_i^ζ are formed from their ML-weighted sums:

$$\frac{x_i^\zeta}{(\sigma_i^\zeta)^2} = \frac{1}{n_c n_r} \sum_{l=1}^{n_c} \frac{x_{n_c(i-1)+l}^\epsilon}{(\sigma_{n_c(i-1)+l}^\epsilon)^2}, \quad (6.8)$$

where only every n_c^{th} rebinned power excess is calculated and n_r is the number of rebinned bins that contribute to each grand spectrum bin (Step 8). The rebinned variances are the analogously scaled, ML-estimated variances given by their combined spectrum counterparts:

$$(\sigma_i^\zeta)^2 = \frac{1}{(n_c n_r)^2} \left(\sum_{l=1}^{n_c} \frac{1}{(\sigma_{n_c(i-1)+l}^\epsilon)^2} \right)^{-1}, \quad (6.9)$$

The rebinned bin frequencies are taken as the average of the contributing combined bin frequencies. For the rebinned spectrum and the subsequent grand spectrum, i now indexes frequency in the more sparsely populated space.

- (8) In order to construct the grand spectrum, which contains both information about the haloscope sensitivity (Step 5) and the axion velocity distribution, the axion lineshape, Eq. (2.14), is calculated given an axion with rest frequency ν_a . The axion lineshape is the probability density function (PDF) of axion particle energies given their Maxwellian velocity distribution in the galactic rest frame, and our planet's velocity through that rest frame (Sec. 2.4.3). This PDF is then discretized into a probability mass function (PMF)¹¹² $\bar{p}_a(\nu_j|\nu_{a,i})$ by integration over rebinned spectrum bins.¹¹³

The grand spectrum is then constructed. The grand spectrum has the same frequency spacing as the rebinned spectrum, but each grand spectrum bin i includes contributions from the

$$n_r = \frac{\Delta_a}{n_c \Delta_b} \quad (6.10)$$

rebinning spectrum bins beginning at i . Sets of $2n_r + 1$ grand spectrum bins hereby become correlated.¹¹⁴ Specifically, the grand spectrum power excesses x_i^η are constructed from the sets of n_r adjacent rebinned power excesses as

$$\frac{x_i^\eta}{(\sigma_i^\eta)^2} = \sum_{l=1}^{n_r} \frac{\bar{p}_a(\nu_{i+l-1}|\nu_{a,i}) x_{i+l-1}^\zeta}{(\sigma_{i+l-1}^\zeta)^2}, \quad (6.11)$$

where ν_{i+l-1} denote the rebinned frequencies. Equation (6.11) provides the ML estimate of the power delivered by an axion to the i^{th} grand spectrum bin, which has variance given by

$$(\sigma_i^\eta)^2 = \left(\sum_{l=1}^{n_r} \left[\frac{\bar{p}_a(\nu_{i+l-1}|\nu_{a,i})}{\sigma_{i+l-1}^\zeta} \right]^2 \right)^{-1}. \quad (6.12)$$

- (9) If the combined spectrum bins did not contain correlations from the SG filtering (Step 4), then the division by the independently obtained sensitivities (Step 5) would ensure that x_i^η/σ_i^η is drawn from a standard normal distribution (no axion) or a normal distribution with

¹¹² This PMF is interpretable as the probability of an individual axion particle being contained in bin j given that the axion rest mass-converted frequency, Eq. (2.9) lies in bin i .

¹¹³ This conversion is made difficult by the fact that the axion frequency will in general not align perfectly with any bin frequency, an effect that matters less the smaller the rebinned bin spacing $n_c \Delta_b$ is. Reference [78] accounts for this uncertainty.

¹¹⁴ The implications of these correlations are addressed in Sec. 7.1.4.

unit mean and standard deviation (axion with $g_\gamma = 1$). The negative correlations imprinted by the SG filter (Sec. 7.1.4), however, reduce the grand spectrum standard deviations σ_i^η by a spectrally uniform factor ξ , which can be extracted directly from the data and validated through simulation.

Scaling the grand spectrum standard deviations down to

$$\tilde{\sigma}_i^\eta = \xi \sigma_i^\eta \quad (6.13)$$

makes the corrected grand spectrum excesses

$$x_i^{(1)} = \frac{x_i^\eta}{\tilde{\sigma}_i^\eta} \quad (6.14)$$

standard normal random variables absent an axion. The superscript (1) denotes that Steps 1–9 have applied for the initial scan, and must be repeated for rescans. Equation (6.14) represents the grand spectrum excesses referred to in the main text.

The SG filter has the additional effect of attenuating the visibility of a potentially present axion. The frequency-independent magnitude of this attenuation can be simulated as well, and is captured along with other, frequency-dependent effects, in Eq. (6.16)'s sensitivity parameters $\eta_i^{(1)}$ (for now, it suffices to say that $\eta_i^{(1)}$ quantifies the mean value of $\langle x_i^{(1)} \rangle$ imparted by a KSVZ axion).

- (10) Steps 1–9 are repeated, possibly multiple times, at frequencies where rescans are deemed appropriate. The corrected grand spectrum excesses and sensitivity parameters are labeled $x_i^{(j)}$ and $\eta_i^{(j)}$, respectively.

In Sec. 7.2, we will evaluate the consequences of each data processing step on our ability to sense an axionic signal in the data.

6.2 Axion search analysis considerations

The analysis starts where the data processing of Sec. 6.1 ends: with a grand spectrum of normalized power excesses $x_i^{(1)}$ measured at frequencies ν_i . The superscript (1) denotes data taken

on an initial scan. Bins displaying large excesses on an initial scan are rescanned, potentially multiple times, yielding rescan spectra $x_i^{(j)}$ indexed as $j = 2, \dots, M$, where many frequency bins i will not be measured in rescans.

For bins not containing an axion, acquisitions of order a millionfold times longer than the inverse bin bandwidth yield, via the central limit theorem, normalized power excesses drawn from the standard Gaussian PDF, obtained by setting $\mu = 0$, $\sigma = 1$ in

$$f_x(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(x - \mu)^2}{2\sigma^2}\right]. \quad (6.15)$$

This distribution is shaded green in Fig. 6.1. For a discussion of the consequences of practical departures from the Gaussian idealization, see Sec. 7.1.5.




Figure 6.1: Gaussian probability density functions $f_x(x)$ for a single-bin normalized power excess x owing to noise alone (green) and noise plus-axion signal for a particular value of the axion-photon coupling g_γ and haloscope sensitivity (blue). In threshold-based inference frameworks (Secs. 6.3, 7.3.1, and 7.3.2), an excess power threshold (black dashed line) is predetermined, and the only information analyzed is whether the measurement is above threshold. The light green (blue) region denotes to a true negative (positive), while the dark green (blue) denotes a false positive (negative). In the Bayesian power-measured (BPM) analysis of Sec. 7.1, the single-scan prior update $u(x)$ (red) is simply the ratio of the two distributions. The exponential dependence of u upon x captures the information content of the measurement. Figure reproduced from Ref. [126].

A bin ν_i that falls at the axion frequency ν_a will have its power excess drawn from what we term the axion distribution. There is in fact a family of such distributions parametrized by the axion-photon coupling g_γ , but it will suffice to consider just one. The axion distribution's standard deviation in the limit of weak axion signal power is approximately that of the no-axion distribution

(i.e. 1) but its mean is offset to

$$\mu_{a,i}^{(j)} = g_{\gamma,i}^2 \eta_i^{(j)}. \quad (6.16)$$

The sensitivity parameter $\eta_i^{(j)}$ is the independently calibrated SNR for an axion with coupling strength $g_\gamma = 1$ in the i^{th} bin of the j^{th} spectrum.¹¹⁵ The PDF of the axion distribution for the i^{th} bin (blue in Fig. 6.1) is thus Eq. (6.15) with $\mu = \mu_{a,i}^{(j)}$, $\sigma = 1$.

The purpose of the statistical analysis for an axion haloscope is to convert the corrected¹¹⁶ grand spectrum excesses $x_i^{(j)}$ and sensitivity parameters $\eta_i^{(j)}$ into probabilistic statements distinguishing between the axion and no-axion hypotheses for a given bin or frequency window. The particle physics community does not lack for statistical methods [235], but the challenge is not in coming up with a mathematically correct framework; rather, it is in finding one which deals well with the particulars of axion searches.

6.2.1 Distinguishing features of axion and haloscope searches

A given analysis framework can be evaluated in large part by how well it accommodates each of the key, defining features of axion searches. Below is a list of some of the most prominent distinguishing traits of axion — and axion-like particle (ALP; see Ch. 9) — dark matter searches. This list is not meant to be comprehensive, nor to perfectly describe every such dark matter search. Each item includes reference to the section of Ch. 7 which addresses it for the statistical framework presented there.

- (1) There is (by assumption) at most one frequency bin containing the axion rest mass among a sea of many bins not containing it. Of interest to the experimentalist is both the probability that the axion lies in any given bin, and in any range of bins. Evaluating the single- and multi-bin hypotheses is covered in Secs. 7.1.1 and 7.1.2, respectively.

¹¹⁵ The parameters $\eta_i^{(j)}$ account for the scaling of the haloscope sensitivity with integration time, experimental parameters, and the dark matter density ρ_a . By normalizing the SNR to $g_\gamma = 1$ and treating the multiplicative coupling factor g_γ^2 as a parameter of the axion distribution, we are following the convention in the axion search literature of assuming a fixed value of ρ_a and setting limits on g_γ (Sec. 2.4.2). The i index on the coupling factor indicates that we need not consider the same axion distribution in each bin.

¹¹⁶ Henceforth we will drop the word “corrected” for brevity.

- (2) Close to the bin containing the axion, there may be bins containing axion-induced power excesses. These and other correlations are handled in Sec. 7.1.4.
- (3) The axion is a persistent signal¹¹⁷ with lineshape that is assumed to be known, at least approximately (Sec. 2.4.3). The persistent availability of the signal contrasts sharply with some other prominent searches for new physical phenomena. For example, in gravitational wave searches, a given astrophysical event [236, 237] must be observed during its one and only occurrence: both false negatives and positives carry important negative consequences. The persistence of the signal is leveraged in all axion haloscope searches through the incorporation of rescan data. See Item 8.
- (4) It is *a priori* unlikely that the axion will reside within the window (mass and coupling) in parameter space to which a given scan is sensitive. It is more unlikely still that the axion will reside within any given bin. These small *a priori* probabilities are dealt with in Sec. 7.1.3.
- (5) Since the axion could (at least in principle) reside over many mass decades (Sec. 2.4.1), linearly spaced Fourier domain bins do not on their own reflect the unstated belief which clearly underlies the current fleet of axion direct detection experiments: that the axion detection community’s belief in axions is nonuniform over frequency. Spectral weighting of the performed hypothesis tests is covered in Sec. 7.1.3. Though the HAYSTAC results of Chs. 7 and 8 are too spectrally narrow for this effect to be of much consequence, the ALP search in Ch. 9 relies heavily upon these considerations.
- (6) The hypotheses for there not being an axion and for there being an axion of any given coupling are generally *both* quantifiable as probability distributions in terms of the measured data. The quantifiability of all hypotheses of interest recommends Bayesian statistics, or at least a reformulation of the null hypothesis that takes into account both distributions.¹¹⁸

¹¹⁷ As mentioned in Footnote 19, sufficiently narrowband axions may actually move noticeably (in the context of a haloscope search) in frequency due to the orbit and rotation of the Earth.

The axion and no-axion hypotheses are directly incorporated in the formulation of Sec. 7.1.1.

- (7) To the extent that the previous item is not true, it is typically primarily because of excess noise power coupled into the detector at discrete frequencies (or ranges thereof). Practical departures from the idealized behavior of the axion and no-axion hypotheses are covered in Sec. 7.1.5.

Additionally, some features more specific to haloscopes and other tuned, narrowband searches (these do not apply, for example, to the ALP search in Ch. 9) in particular are:

- (8) Frequency bins are scanned in groups that are small compared to the total number of bins scanned during an entire data run.¹¹⁹ As such, targeted, narrowband rescans can be performed with relative ease on power excesses that appear in an initial scan. Section 7.3 showcases key advantages of the framework in Ch. 7 pertaining to how rescan data may be incorporated.
- (9) For a QCD axion (assuming it does exist, and has the characteristics put forth in Ch. 2) singled out by a fully operational haloscope, nature has designed the expected signal behavior to be almost fool proof.¹²⁰ Specifically, its particular scaling with applied magnetic field and the cavity transmission profile, together with its lineshape and persistence, makes it hard for spurious RF excesses to masquerade as axionic. This last point creates a fundamental asymmetry between exclusion and discovery. It has important implications for how false positives and negatives should be guarded against in the analysis. Those implications are discussed in the context of accounting for the look-elsewhere effect, Secs. 6.5 and 7.1.2.

¹¹⁸ Sharon Bertsch McGrayne traces the history of the development of Bayesian statistics in Ref. [238]. A quote from her work is pertinent to our discussion: “Egon Pearson’s idea was that the only correct reason for rejecting a statistical hypothesis was to accept a more probable one.”

¹¹⁹ While broadband searches are possible [110], it has been shown that a narrowband/resonant (quality factor Q) approach is a Q -fold more effective strategy in terms of scan rate than using a broadband receiver [95].

¹²⁰ However, as Edward Teller purportedly once said, “if you think you have made something fool proof, the fool is always bigger than the proof” [239].

6.3 Frequentist threshold framework

The majority of existing axion exclusion has been obtained using either of two frequentist methods. The haloscope exclusion literature [78, 89, 91, 93, 94, 97–109, 115] often does not do a good job distinguishing which framework is being used in which publication. As will be made clear below, the two methods mean substantively different things by “exclusion,” so this ambiguity is a rather unfortunate one. The first of the two, described in this section, uses a (typically) multi-step threshold-based test to either exclude the axion or perform more advanced tests on a persistent signal. The purpose of this section is twofold: first, along with the next section, it aims to lay out for readers exactly how frequentist exclusion is and has been obtained for haloscopes. Second, the framework in this section can be compared directly to one presented in the next chapter. The remainder of this chapter, together with the next, constitute an argument that the frequentist threshold framework considered in this section is on balance better suited to the task of axion haloscope detection than that in Sec. 6.4, yet worse than that in Ch. 7.

6.3.1 Framework details

In order to perform an axion search subject to the FT framework, a power excess threshold $x_T^{(j)}$ is set for each scan j and this threshold is used to determine firstly whether or not to proceed with rescans, and secondly whether or not to exclude the axion hypothesis at each frequency. The logic of this exclusion is prescribed as follows. First, the experimentalist establishes a null hypothesis. In this case, the null hypothesis for each bin is that there *is* an axion at the bin’s frequency.¹²¹ Second, the experimentalist establishes and follows a procedure for acquiring data. This procedure may include conditional steps: e.g. “acquire more data in a given bin via a rescan if and only if the initial scan power exceeds the threshold set for that bin.” All possible paths of action eventually terminate either in a rejection of the null or a failure to reject the null, forming a decision tree. Figure 6.2 shows how a realistic grand spectrum of frequency bins might be scanned

¹²¹ The axion hypothesis is chosen as the null so that it might be rejected, or excluded, later. Because of the unconventional choice to make the interesting (i.e. axion) hypothesis the null, the language of false negatives and positives will adopt the unconventional usage where a positive (negative) refers to the null being true (false).




Figure 6.2: Illustrative, fictitious grand spectrum data for a standard haloscope subject to a frequentist threshold (FT) analysis allowing $M = 4$ scans. The initial scan data from a typical haloscope [105, 107] consists of a normalized power excess in each of $N \sim 10^5\text{--}10^6$ frequency bins, represented above by the boxes in the first row. At frequencies where the excess is above a predetermined threshold (red boxes), additional scans are conducted until a measurement below threshold is recorded. Only if all M scans exceed threshold is failure to reject the axion hypothesis reported. If any scan comes in below threshold (black boxes), the axion hypothesis is rejected for the bin in accordance with Eq. (6.20) and further scanning need not be conducted (empty boxes). For the FT framework of Sec. 6.3, both the threshold and the number of scans M must be predetermined, and the exclusion therefore may take into account scans that were not performed. Here, Scan 4 was not performed for any bins, but still impacts the final reported exclusion. Because the haloscope bandwidth is typically larger than the expected axion linewidth, bins adjacent to those exceeding threshold are automatically rescanned (gray boxes). However, the FT framework discards these data regardless of the power measured. Conversely, the BPM method discussed in Ch. 7 is able to use the information. Figure reproduced from Ref. [126].

according to the powers measured relative to threshold. Exclusion is set for all bins i where the null is rejected at a confidence level of

$$E = 1 - F_n, \quad (6.17)$$

where F_n is the bin- i false-negative rate of the entire procedure: the total probability that the null would have been rejected, were it true.¹²² We refer to this class of methods as frequentist threshold

¹²² This approach is equivalent to defining a test statistic

$$t_s = \begin{cases} 1 & x^{(j)} \geq x_T^{(j)} \forall j \\ F_n(g_\gamma) & \text{otherwise} \end{cases} \quad (6.18)$$

from the total false-negative rate treated as an explicit function of g_γ . Couplings where t_s goes below some predetermined target rate $F_{n,\text{target}}$ are rejected at $1 - F_{n,\text{target}}$ confidence level via Eq. (6.17). The piecewise definition of t_s , Eq. (6.18), yields exclusion only when at least one scan's power excess $x^{(j)}$ comes in below its threshold $x_T^{(j)}$. When all scans exceed threshold, $t_s = 1$ indicates the presence of (possibly axionic) excess power.

(FT) frameworks. FT frameworks answer the question, “assuming the axion exists, what is the probability of failing to observe it?”

An FT search procedure [107] is prepared to perform as many as M scans at each frequency ν_i , reporting exclusion subject to its predetermined false-negative rate. At a given ν_i , no additional scans are performed once any scan’s measured power excess $x_i^{(j)}$ fails to exceed its predetermined threshold $x_T^{(j)}$, the vertical, dashed line in Fig. 6.1. For the general case where each scan may use a different false-negative rate (dark blue region of Fig. 6.1), given from integrating Eq. (6.15) by

$$f_n^{(j)} = \int_{-\infty}^{x_T^{(j)}} f_x(x; \mu_a^{(j)}, 1) dx, \quad (6.19)$$

the total false-negative rate is

$$F_n = \sum_{j=1}^M f_n^{(j)} \prod_{k=1}^{j-1} (1 - f_n^{(k)}). \quad (6.20)$$

This total false-negative rate is what sets the confidence level of the reported exclusion E as in Eq. (6.17). For the special case where all M scans have equal false-negative rates, $f_n^{(j)} = f_{n0}$, Eq. (6.20) reduces to

$$F_{n0} = 1 - (1 - f_{n0})^M. \quad (6.21)$$

In practice, the false-negative rates $f_n^{(j)}$ and power thresholds $x_T^{(j)}$ are set to convenient values, removing the frequency-dependence of $\mu_a^{(j)}$ in Eq. (6.16) by forcing g_γ to compensate for the frequency-dependence of the sensitivity. The initial scan sensitivities $\eta_i^{(1)}$ therefore determine the minimum coupling $g_{\gamma,i}$ that is excluded (or not) at each frequency. Subsequent ($j > 1$) scans, conversely, integrate to a sensitivity $\eta_i^{(j)}$ determined by the coupling $g_{\gamma,i}$ reached on the initial scan. The result, in the case where the axion hypothesis is rejected for all bins where it can be tested, is an exclusion plot [107] whose spectral structure reflects that of $\eta_i^{(1)}$.

The procedure for collecting data and reporting exclusion in these FT frameworks must be rigidly defined and laid out in advance. If, instead, the experimentalist is allowed to alter the decision tree that leads to negative or positive results while navigating that decision tree, the result becomes susceptible to bias. To this point, the total false-negative rate F_n , Eq. (6.20), directly responsible for setting the exclusion E , grows with the total number M of scans in the predetermined

experimental procedure. Suppose, then, that the experimentalist was willing to perform as many as M scans at each frequency bin i before failing to reject the axion hypothesis, but that in practice some number $m < M$ was all that was required to get at least one negative result in every bin. The correct exclusion to report, perhaps unintuitively, is that which takes into account all M allowed scans, including even the $M - m$ unperformed scans. The number of such unperformed scans and their individual false-negative rates $f_n^{(j)}$, Eq. (6.19), must be known to properly report exclusion. If not written down in advance, these numbers are difficult to estimate without inserting bias. In practice it is extremely difficult to rigorously adhere to FT logic. For example, one would expect that some exclusion would be reported for a total number of scans M surpassing m , the exact number performed, yet this never appears in the literature [89,94,97–104,106–108].¹²³ Furthermore, exclusion has been reported for a number of scans less than the number performed [78]!

Though the preceding logic and inference is frequentist in nature, an axion haloscope analysis can alternatively be carried out using the language of Bayesian statistics. In Sec. 7.3.2, I present a threshold-based Bayesian analysis (denoted BT2) which is equivalent to the FT framework, given simple assumptions applicable to all dark matter axion experiments. That is, the two frameworks share operational procedures, and they ultimately output identical conditional probabilities. This correspondence implies that we can quantitatively interpret FT axion exclusion as decreased belief in, or probability of, the axion hypothesis.

6.3.2 Aggregate exclusion

The aggregate exclusion $\mathcal{E}(g_\gamma)$ will serve as our means of applying the logic of the FT framework to an entire scanned window (or portion thereof). We will see that $\mathcal{E}(g_\gamma)$ provides a more conservative, yet more faithful assessment of the exclusion achieved by a haloscope than does the standard practice of averaging the couplings at which a given exclusion E is achieved.

The aggregate exclusion $\mathcal{E}(g_\gamma)$ answers the same question that the exclusion $E(g_\gamma)$ (written

¹²³ The references cited here are those which either definitely use the FT framework or might plausibly do so. A few publications which are not clear on framework usage have been determined to be using the framework of Sec. 6.4 with specific modifications discussed therein via private correspondences with the authors.

explicitly as a function of coupling g_γ) does, but for the entire frequency range. The exclusion $E(g_\gamma)$ answers the question: “supposing an axion of coupling g_γ occupies a given bin, what is the probability of having observed a predefined positive result (i.e. detection of the axion) there?” We therefore define the aggregate exclusion $\mathcal{E}(g_\gamma)$ as the chance of having observed a predefined, global positive result (since it is not known which bin the axion is in, the global positive result is defined to mean at least one positive result across the N frequency bins.), assuming the axion rest mass to occupy exactly one bin within the range. In keeping with the uniform-priors assumption of Sec. 7.1.3, the likelihood of an axion to be at any given bin is equal for all bins, and for simplicity the powers in bins are assumed otherwise uncorrelated (see Sec. 7.1.4). In Sec. 6.3.1, the known sensitivities $\eta_i^{(j)}$ are thought of as combining with predetermined power thresholds $x_T^{(j)}$ to preselect couplings $g_{\gamma,i}$ to test at a uniform, desired false-negative rate F_n . This logic can be partially inverted: assuming the same $\eta_i^{(j)}$ and $x_T^{(j)}$, now use a frequency-independent coupling g_γ to preselect frequency-dependent total false-negative rates $F_{n,i}(g_\gamma)$ at which to test. This framing allows for the testing of a uniform coupling g_γ across the entire range to a confidence level of

$$\mathcal{E}(g_\gamma) \approx \frac{1}{N} \sum_{i=1}^N [1 - F_{n,i}(g_\gamma)], \quad (6.22)$$

where the approximation is valid in the low total false-positive rate limit of Sec. 7.3.2, $F_p \ll 1$. The confidence to test at will increase monotonically with the coupling g_γ being tested. The coupling coincident with a desired percentile (e.g. 90%) can be determined from the sensitivity profile $\eta_i^{(j)}$ alone, and a standard confidence test performed by comparing the measured power excesses $x_i^{(j)}$ against the predetermined thresholds $x_T^{(j)}$, with a positive result defined as a click on all M scans in at least one bin.¹²⁴

The aggregate exclusion, by answering the question of threshold-based frequentist confidence testing, proves a more natural figure of merit for FT exclusion over a given scan window than the standard practice [107] of quoting the average coupling g_γ^{mean} excluded to a given confidence

¹²⁴ The aggregate exclusion $\mathcal{E}(g_\gamma)$ bears exactly the same relationship to the aggregate prior update $\mathcal{U}(g_\gamma)$ of Eq. (7.7) that $E(g_\gamma)$ bears to $U(g_\gamma)$ of Eq. (7.6), making it central to comparing the Bayesian framework of Ch. 7 to the frequentist framework of this section.

E . If the coupling g_γ^{agg} is defined as that for which identical aggregate exclusion $\mathcal{E}(g_\gamma^{\text{agg}}) = E$ is achieved, then a comparison will reliably reveal $g_\gamma^{\text{agg}} > g_\gamma^{\text{mean}}$. The standard metric, g_γ^{mean} , is not somehow achieving deeper sensitivity, it is averaging in a space inconvenient to the natural logic of frequentism, and biased relative to it. More directly, an experiment that achieves 90% exclusion at a frequency-averaged coupling of g_γ^{KSVZ} has *not* rejected the hypothesis that a KSVZ axion lies within its scan window with 90% confidence.

6.4 p -value framework

The second commonly used framework for excluding axions is a standard frequentist method in particle physics [235] that relies on p -values, which I will argue is ill-suited to the particulars of axion direct detection. To distinguish it from the frequentist framework of Sec. 6.3, I will refer to it simply as the “ p -value framework.” The p -value, or probability value, answers the question: “what is the probability that an event at least as extreme as the one observed would have occurred, subject to the null hypothesis?” The event in question is the measurement of a grand spectrum power $x_i^{(j)}$ (or, more generally, a ML-weighted average of such powers over multiple scans). For the p -value method in its simplest form, exclusion at confidence level E_p is reported for the axion g_γ which satisfies

$$F_x \left(x_i^{(1)}; \mu_{a,i}^{(1)}, 1 \right) = 1 - E_p, \quad (6.23)$$

where

$$F_x(x; \mu, \sigma) = \int_{-\infty}^x f_x(x'; \mu, \sigma) dx' \quad (6.24)$$

is the cumulative distribution function (CDF) obtained from integrating Eq. (6.15), here evaluated at the measured power excess $x_i^{(j)}$ with $\mu = \mu_{a,i}^{(1)}$ and $\sigma = 1$. The coupling g_γ enters through $\mu_{a,i}^{(1)}$ via Eq. (6.16).

The coupling excluded at each frequency is therefore the one for which the measured power or a more extreme (that is, lower, for the inverted use of the axion hypothesis as the null hypothesis for frequentist exclusion; Footnote 121) power would have been measured only $1 - E_p$ of the time,

were the axion really at that coupling. Intuitively, this sounds fine. However, p -values are famously misleading quantities. The American Statistical Association in 2016 issued a statement cautioning strongly against the widespread reliance on p -values found throughout the sciences [240]. They are particularly dangerous when many independent hypotheses are being tested, giving rise to the phenomenon of p -hacking. If not used with care in the context of axion direct detection, we will see in Sec. 6.5.1 that this danger applies here as well, in a modified form.

To be sure, there is nothing logically incoherent about a p -value. It is simply a mathematical object; any error is always on the user end. Let us consider two simple pitfalls associated with the use of p -values in the axion exclusion context. We will consider a third pitfall associated with the testing of many independent hypotheses in the next section.

First, a p -value essentially treats the measured power as itself-plus-a-probability-weighted-sum-of-all-more-extreme-values, for the purpose of excluding. While there is nothing logically wrong with this, an elementary, intuitive understanding of (say) 90% frequentist exclusion might be that an axion at the excluded coupling would have been identified 90% of the time, were it there. This is precisely the quantitative meaning of the FT exclusion of Sec. 6.3. However, by grouping the measured power only with that subset of unmeasured powers which are more extreme (i.e. more favorable to more aggressive exclusion) than it, the p -value exclusion reliably rules out axions which would *not* have set off even a single rescan with anything close to 90% probability, were they present.

This pitfall is greatly exacerbated by a second feature of the p -value method: that it still often uses a threshold to test for discovery, and that that threshold takes into account the look-elsewhere effect only with respect to discovery, yet not exclusion (Sec. 6.5). To give an example, a single-scan threshold may be set at 3.5σ (equivalently, $x_T^{(1)} = 3.5$). A typical bin will come in around 0σ , causing the axion at 1.28σ to be excluded at $E_p = 90\%$ confidence level via Eq. (6.23). The probability that such an axion (mean at 1.28σ) would have actually set off the detector for even a single rescan is, from Eq. (6.19), 1.3%. This can be made worse if, as is common, rescans are performed. That is to say, when the FT framework reports 90% exclusion at a given g_γ , it means

that an axion would have been identified 90% of the time. Indeed, we saw in Sec. 6.3.1 that the FT framework takes care to ensure that the exclusion is the complement of the false-negative rate for an arbitrary M -scan protocol. But when the p -value framework reports 90% exclusion for an axion a given g_γ , the most typical result is that such an axion would have been identified 1.3% of the time for the initial scan, and subsequent scans may add additional opportunities to miss the signal.¹²⁵ What would have in fact typically happened if an axion existed at 1.28σ is that an axion at a typical value of 2.56σ , still well short of 3.5σ , would have been excluded. In a context where of order 100,000 bins might be tested in a single scan, and there are not nearly that many pixels displayed in a typical exclusion plot, this result could very well do nothing to modify even a single pixel of the exclusion as displayed. In spite of all of this, there is nothing logically wrong with the p -value framework: it is simply that most readers are not conditioned to look at an exclusion line and understand that it might imply a typical false-negative rate of nearly 99%.

The second pitfall of the p -value framework is that it tends to exclude nonsensical couplings if left unaltered. Simply put, any measured power below -1.28σ will require excluding an axion with mean below 0σ at 90% confidence level. Of course, there is no axion with subzero mean, so all axions are excluded. This is true independent of sensitivity. So in principle, an axion search could be carried out not in the typical ~ 8 T fields of modern haloscopes, but in the Earth's field of ~ 50 μ T, and 10% of the (say) 100,000 frequency bins would be excluded at *all* couplings down to $g_\gamma = 0$. Again, there is nothing logically wrong here. The 10% of bins that achieve this nonsensical result of perfect exclusion are what the quoted confidence level of 90% warns about, but the result is highly inconvenient. One remedy is to artificially modify the axion distributions in a manner such that these results never happen, but it seems a strange price to pay when alternative frameworks exist which allow for the use of the actual axion PDF.

The third pitfall, which we will see in the next section, is that the exclusion reported over the

¹²⁵ This story is complicated by the fact that adjacent bins tend to be scanned during rescans (Fig. 6.2). The possibility that a bin with an axion might have been missed on the first scan, but showed up on the rescan of a neighboring bin which did exceed threshold, can partially mitigate the extremely high effective false-negative rate. Both the exacerbating and mitigating effects of rescans depends on the particulars of the rescan criteria and parameters, as well as the strength of correlations in nearby bins.

entire frequency range for any properly functioning framework ought to get reliably worse as more independent frequencies are probed. The p -value method or close variants of it [75] as typically implemented in axion searches [91, 93, 105, 109, 115] appear to not fully account for this effect, whereas the other frameworks discussed in this thesis all do. For all of these reasons, the p -value method should only be used with expert handling, or better yet not at all, given the availability of more intuitive alternatives (Secs. 6.3, 7.3.1, 7.3.2, and especially 7.1).

6.5 Look-elsewhere effect: implications for detection and exclusion

The look-elsewhere effect is the name for the phenomenon whereby rare events become commonplace when many independent hypothesis tests are performed. It must commonly be dealt with in particle physics searches [235] which entail testing many independent parameter values simultaneously. Typically, accounting for the look-elsewhere effect means shifting from a notion of how locally unlikely an event is to how globally unlikely it is. Here, local refers to what occurs in any given bin, and global refers to what happens in all the bins taken as a collective.¹²⁶ For example, an event that has a 0.01% chance of occurring locally, for example, will occur globally roughly 1% of the time that 100 independent tests are performed.¹²⁷ In a rare event search, therefore, a 5σ global result is a typical scientific standard for discovery. However, in an axion haloscope context, accounting for the look-elsewhere effect is the conventional sense — that is, with respect to discovery (a modifier that will help us is distinguishing from its implications for exclusion, discussed below) — is somewhat trivial. In principle, a total unawareness that rare events occur more frequently over many trials could lead a naïve experimentalist to set a much-too-low rescan threshold and therefore have to perform an over-abundance of rescans; in practice, that mistake does not appear to be being made in narrowband, tuned experiments. Conversely, in broadband experiments without targeted rescan capabilities, it is prudent to explicitly account for the look-elsewhere effect when

¹²⁶ The quantities in this thesis defined with “aggregate” in the name, Eqs. (6.22) and (7.7), answer questions about global events.

¹²⁷ A global occurrence is here taken to mean one or more local occurrence over a set of tests. This definition is harmonious with the hypothesis that an axion may lie in at most one bin, Item 1 in Sec. 6.2.1.

reporting discovery [75, 241]. Even in such contexts, however, the availability of simple and robust manual tests can render the notion of reporting discovery based on a limited number of identical scans overly simplistic. In searches for hidden photon dark matter [242] where the magnetic field plays no role [243, 244], accounting for the look-elsewhere effect with respect to discovery is also important, and doubly so if those searches are also broadband.¹²⁸ This might lead us to believe that accounting for the look-elsewhere effect is unimportant for axion haloscope searches. In fact, the opposite is true.

In the context of an axion search, haloscopic or otherwise, the look-elsewhere effect worsens the exclusion that is achievable across a scanned frequency window. The coupling excluded will move upwards¹²⁹ as more independent tests are conducted, reflective of the increasing difficulty in ruling out a special frequency when it has more imposters to hide among.¹³⁰ This effect is not fully accounted for in the exclusion of some axion searches [75, 91, 105, 111]. If accounted for, it would have the effect of scaling back existing exclusion results.¹³¹

We can be quantitative about the role the look-elsewhere effect plays in slowing down a scan. Scan rate is traditionally calculated [42] as in Eq. (2.16), without the $\text{SNR}^2 \rightarrow \text{SNR}'^2$ correction to the denominator of the first term. As promised in Sec. 2.5.2, we will here estimate that correction term. This is not a rigorous derivation, but just a quantitative means of seeing what the look-elsewhere effect very roughly does to scan rate.

Without any involved calculations of the sort performed in Sec. 7.3.3, we can anticipate that a very rough criterion to be sensitive to an axion of some coupling g_γ across an entire scan is that

¹²⁸ Another, less common example, but one relevant to this thesis (Sec. 7.4), is a haloscope that takes some data and promptly stops working, because of, for example, a magnet quench [42, 90].

¹²⁹ This effect can be observed in the numerical simulations of Sec. 7.3.3.

¹³⁰ This fact is analogous to flipping N coins M times apiece and looking to rule out the existence of one special, anonymous coin with heads probability $p_h > 0.5$ hidden among the whole set. The more coins you flip, the weaker the overall p_h you should be able to rule out. It is a question of semantics whether to call this the look-elsewhere effect, which could be thought of as only applying to discovery by construction. This thesis takes the perspective that because the effects on discovery and exclusion both originate directly from the tendency of rare events to occur more when there are multiple trials, it is appropriate to view them both as manifestations of the look-elsewhere effect.

¹³¹ More accurately, there is nothing wrong with the existing exclusion results from a purely logical perspective. Rather, it is that, overlooked for sufficiently many independent tests, the logic itself ceases to imply that (frequentist) one would have detected the axion, were it in the scan window, or that (Bayesian) the probability of its being there actually went down in light of the data. Subject to any reasonable definition of “exclusion,” this functionally means that the reported exclusion should be scaled back (i.e. upwards in coupling).

the expected number of background fluctuations $\langle N_B \rangle$ that look like that axion obeys

$$\langle N_B \rangle \ll 1. \quad (6.25)$$

This is because there can (by assumption) be at most one axion, and in any realistic scan of $N_I \gg 1$ independent bins,¹³² N_B will be Poisson-distributed, with standard deviation equal to the square root of its mean

$$\sigma_{N_B} = \sqrt{\langle N_B \rangle}. \quad (6.26)$$

Therefore, once $\langle N_B \rangle$ gets close to unity, so does σ_{N_B} , and it becomes impossible to reliably tell a whether a single additional axion-like excess is due to an axion, or due to the fluctuations expected from the noise statistics.

To a rough approximation, what it means for a fluctuation to be axion-like (i.e. to count towards N_B) is for it to achieve a grand spectrum power excess x of at least some significant fraction of the SNR of the axion distribution.¹³³ To put a number on it for the purposes of this calculation, we will say that a fluctuation is axion-like if it has¹³⁴

$$x \geq \text{SNR}. \quad (6.27)$$

Therefore,

$$\langle N_B \rangle = N_I [1 - F_x(\text{SNR}; 0, 1)], \quad (6.28)$$

where $F_x(\text{SNR}; 0, 1)$ is the no-axion CDF, Eq. (6.24), here evaluated at the SNR per our criterion, Eq. (6.27). To proceed further, we must understand the role of the SNR as it appears in the equation for scan rate, Eq. (2.16).

When Eq. (2.16) is written in its uncorrected form (with SNR and not SNR'), a high SNR (e.g. of around 5, perhaps inspired by the common 5σ standard for discovery) is usually assumed. However, as discussed above, in haloscope searches it is not necessary to build one's fortifications

¹³² We will here think of the number of bins as some effective number N_I of independent bins, to keep the math simple.

¹³³ in Sec. 6.3, we labeled this scan- and bin-dependent quantity $\mu_i^{(j)}$, here we simply recognize it as the SNR, since we are not working at a level of detail where we care about scan- or bin-dependencies.

¹³⁴ One could argue for the use of a multiplicative constant in this equation, or perhaps even for a different criterion entirely. The point here is largely just to show in what form N_I enters the scan rate.

against accidental discovery into the statistical framework alone, when there are robust tests with vanishingly low error rates that can be performed on persistent signals. Rather, the high values of the SNR usually chosen are, consciously or not, set to avoid paying the overhead of performing these time-expensive tests in a context where many independent tests are performed. In other words, the look-elsewhere effect with respect to exclusion is already to zeroth order built into most estimates of the scan rate by choosing a value for the SNR that only makes sense in a global context. But either for a region of interest much smaller or (more practically, when considering the ambitions of the field as a whole) larger than a typical scan, we can anticipate how the SNR term should change to accommodate the different number of bins, while obeying Eq. (6.25).

Subject to the logic of the previous paragraph, a cleaner interpretation of SNR and SNR' is to temporarily think of SNR as the mean axion power one would desire if one were just scanning (unrealistically) a single bin. The desired value will depend on the overhead cost of performing time-expensive tests off cavity resonance or with the magnet ramped down, or, if in a situation (perhaps an ALP search) where such tests are less fool proof, the (much scarier) risk of accidentally reporting a false discovery. In any case, if we now go from 1 to N_I independent bins, we may ask how much higher the local SNR, which we now relabel SNR', must get in order to keep $\langle N_B \rangle$ constant: that is, how far we have to back off in coupling to maintain our global SNR given a fixed scan time.

To compare the two cases ($\{1 \text{ bin, SNR}\}$ versus $\{N_I \text{ bins, SNR}'\}$), where

$$\text{SNR}' = \beta(N_I) \text{SNR} \tag{6.29}$$

is parametrized by a quantity $\beta(N_I) \geq 1$ to be determined, we write out Eq. (6.28) for each:

$$\langle N_B \rangle_{\text{case1}} = \frac{1}{2} \text{erfc} \left(\frac{\text{SNR}}{\sqrt{2}} \right) \tag{6.30}$$

$$\langle N_B \rangle_{\text{case2}} = \frac{N_I}{2} \text{erfc} \left(\frac{\beta \text{SNR}}{\sqrt{2}} \right), \tag{6.31}$$

where erfc is the complimentary error function. Equating Eqs. (6.30) and (6.31) to preserve the

expected number of background, axion-like fluctuations, we find

$$\beta(N_I) = \frac{\sqrt{2}}{\text{SNR}} \text{erfc}^{-1} \left[\frac{1}{N_I} \text{erfc} \left(\frac{\text{SNR}}{\sqrt{2}} \right) \right], \quad (6.32)$$

where erfc^{-1} is the inverse complimentary error function. According to Eq. (6.29),

$$\begin{aligned} \text{SNR}' &\sim \sqrt{2} \text{erfc}^{-1} \left[\frac{1}{N_I} \text{erfc} \left(\frac{\text{SNR}}{\sqrt{2}} \right) \right] \\ &\approx \sqrt{\text{SNR}^2 + 2 \log N_I}, \end{aligned} \quad (6.33)$$

where the last approximation, equivalent to Eq. (2.17), is seen in Fig. 6.3 to be valid at low N_I and high SNR. The \sim sign is a reminder that this is only a rough calculation of the effect of N_I . When $N_I \rightarrow 1$, $\text{SNR}' \rightarrow \text{SNR}$, as expected.




Figure 6.3: Effective correction to the required SNR from the look-elsewhere effect. Curves for different global values of SNR showing how the local value SNR' must increase to accommodate higher numbers N_I of independent bins and maintain sensitivity to a fixed axion. The solid curves correspond to the exact form (first line) of Eq. (6.33) and the dashed curves to the approximate form (second line).

The interpretation of Eq. (6.33) is that going to a higher number of independent frequency bins N_I entails searching for more strongly coupled (equivalently, higher local SNR') axions. Figure

6.3 shows the effective look-elsewhere correction to the local SNR' for several different global values of SNR. For example, for a scan like Phase I, Run I of HAYSTAC [78, 106] containing of order 10^5 independent bins and integrating to a local SNR' of ≈ 5 ,¹³⁵ we first determine that the global $\text{SNR} = 2$ curve intersects $\text{SNR}' = 5$ at $N_I \approx 10^5$. The same curve passes through 10^6 around $\text{SNR}' \approx 5.5$. Therefore, for a HAYSTAC Phase I, Run I-like scan to exclude the same couplings over a scan of $10\times$ as many bins, it would take an additional $(5.5/5)^2 - 1 \approx 20\%$ longer according to Eq. (2.16) than naively expected from the tenfold increase in number of bins alone. This is a somewhat reassuring result: it means that for most scans of practical sizes, scaling up even by an order of magnitude does not incur much further penalty. However, a $\sim 20\%$ effect is not so small as to be totally negligible. And, as we saw in Sec. 6.4 and will discuss further momentarily, when frameworks do not account for the look-elsewhere effect with respect to exclusion at all, the effect is considerably more dramatic.

6.5.1 Frequentist frameworks and the look-elsewhere effect

Now we can ask, do the frequentist frameworks discussed in this chapter account for the look-elsewhere effect? The FT framework does, albeit in a somewhat clunky way. In order to avoid an over-abundance of false positives, the threshold $x_T^{(j)}$ must be moved up as more and more bins are tested. Since the same axion that is tested for rescanning (in the case of a multi-scan protocol) is also tested for exclusion, this necessarily implies testing more strongly coupled axions at fixed confidence level. Hence, while the choice of exactly what (local) false-positive rate to set is left somewhat subjectively up to the experimentalist, that rate must eventually get lower as more and more bins are tested. The result is a worse exclusion line and worse aggregate exclusion, Eq. (6.22), as required by the look-elsewhere effect.

The p -value framework in its simplest form (Sec. 6.4) accounts for the look-elsewhere effect only with respect to discovery. The threshold used to test for rescans (in a multi-scan protocol) has nothing directly to do with what coupling is excluded, as discussed in Sec. 6.4. While this

¹³⁵ The actual value used was 5.1.

threshold moves up for the same reason that the threshold in the FT framework does as more bins are scanned, the couplings excluded do not move up with it. The basic version of the p -value framework therefore does not account for look-elsewhere with respect to exclusion, which explains why the reported confidence level largely loses its intuitive meaning.¹³⁶

The look-elsewhere effect can be accounted for in the p -value method numerically. Over a given subset of N_s frequency bins, the coupling that would have produced the highest measured power or more extreme 90% of the time (assuming all the other bins had no axion, and still generate powers from the no-axion distribution), can be excluded over the subset of bins at 90% confidence. If that subset is made identical with the entire set of bins for the data run,¹³⁷ the resulting exclusion properly accounts for the look-elsewhere effect with respect to exclusion: that is, it typically gets worse as more bins are probed. In practice, this does not appear to be being done for exclusion using the p -value framework. Though in at least some cases [91, 105] it is being done in some form for subsets of bins at the pixel scale (i.e. over however many bins fit into a single frequency-axis pixel in the exclusion plot as displayed in a publication). There, the look-elsewhere effect is being accounted for properly at that spectral scale, but not beyond it. To the degree that the exclusion plot consists of more than one independent pixel, the reported exclusion becomes progressively more untethered from the notion that an axion, if present, would have been observed. In short, the physics of the early universe did not know about the size of a printed page on which results about it would one day be displayed.

In the next chapter, we will discuss a Bayesian analysis framework that is well suited to deal with all of the distinguishing features of axion searches listed in Sec. 6.2.1, including the look-elsewhere effect with respect to both discovery and exclusion.

¹³⁶ This is closely related, but not identical, to the first p -value pitfall discussed in Sec. 6.4. In particular, the notion that the exclusion should be getting worse is an alternate framing of the exacerbation of the first pitfall whereby, if it doesn't, the reported exclusion loses its intuitive meaning. However, the p -value's effective grouping of the measured power with only the subset of unmeasured powers which are more extreme than it is a distinct and separable facet of the problem. Both facets conspire to produce the particularly egregious 98.7% effective false-negative rate in the example in Sec. 6.4.

¹³⁷ In fact, there is no reason other than convenience to box in just the data run in this fashion. The look-elsewhere effect with respect to exclusion should properly be accounted for *between* different data runs and experiments.

Chapter 7

A maximally informative axion dark matter analysis

In Ch. 5, we saw how quantum measurement may be used to speed up axion searches, albeit at a tangible hardware and complexity cost to the experiments. While the ultimate means of bringing total haloscope scan times down from the formidable (millennia-plus) timescales of Sec. 2.5.3 will have to involve hardware advances, the lowest hanging fruit for improving axion detection lies within the data processing and analysis. Subject to the distinguishing features of axion searches (Sec. 6.2.1), a statistical framework can be more or less efficient with the information content of the measurement, and allow for more or less operational freedom on the part of the experimentalist. This chapter, together with the preceding one, presents the work of Ref. [126]. In it, I argue that an optimization of the statistical analyses that haloscopes have historically used is readily available. The standard threshold-based confidence tests used in the frequentist threshold (FT) analysis of Sec. 6.3 transform continuous measurements of power into binary outcomes, discarding knowledge of what power was measured, and, with it, sensitivity to the axion. Because haloscopes are in practice statistics-limited experiments, for which additional data continues to enhance sensitivity, a more informative analysis will translate into tangible time savings during operation.

In Sec. 7.1, I will introduce the Bayesian power-measured (BPM) framework, a statistical analysis framework for axion searches that makes intuitive use of all of the information at the output of the data processing. We will see in Sec. 7.2 that the processing itself is nearly information-optimal. Other sections exist to show how BPM accommodates the unique features of axion searches put forth in Sec. 6.2.1 and to provide useful insight into BPM through comparison with competing

frameworks. Section 7.3 shows that BPM permits operational freedoms not available in its strongest competitors. These freedoms can be leveraged in future work to determine more optimal scan strategy [95], the other hardware-free advance which to date still has not been rigorously explored. The BPM framework is put to the test in Sec. 7.4 by reanalyzing the HAYSTAC Phase I dataset [107]. The framework provides improved constraints on the axion-photon coupling g_γ while also identifying the most promising regions of parameter space within the 23.15–24.0 μeV mass range. A comparison with the standard threshold analysis suggests a 36% improvement in scan rate from this analysis. In addition, it spotlights within the dataset the group of frequencies “least unlikely” to contain the axion. The most prominent among these, though still unlikely to contain an axion, nonetheless stands out relative to any other frequency within the scan range.

7.1 Bayesian power-measured framework

This section introduces the BPM analysis framework, which simplifies the operational constraints placed on experimentalists while making use of all of the information content of coherent axion detection data. This framework builds on the existing HAYSTAC processing procedure [78] and may be readily adapted to other experiments.

Adopting a Bayesian perspective, we can ask whether a more informative analysis than the FT framework of Sec. 6.3 is possible, without the downsides of p -values (Sec. 6.4). We consider frameworks which operate by applying Bayes’ theorem,

$$P(Y|Z) = \frac{P(Z|Y)P(Y)}{P(Z)}, \quad (7.1)$$

to an axion search data set. Notationally, $P(B)$ denotes the probability of event B being true, while $P(B|C)$ denotes the same, conditional upon event C being true. In Eq. (7.1), $P(Y)$ is the Bayesian prior probability of event Y being true, and it is updated by the occurrence of event Z to the posterior probability $P(Y|Z)$. Section 7.3 describes two distinct Bayesian threshold (BT) frameworks in which Z is taken to be a binary outcome, or set thereof. We will see there that a more informative analysis is indeed possible, and that thresholding always imposes some operational

restrictions and needlessly discards valuable information. Nonetheless, the BT logic exemplifies how Bayes’ theorem may be applied to axion detection. In doing so, it builds an important bridge between the FT and BPM frameworks.

7.1.1 Updated belief

Bayes’ theorem, Eq. (7.1), may be applied to haloscope data so as to preserve the full information content of the measurement. The information at each frequency may then be aggregated (Sec. 7.1.2) into a statement about the change in probability of an axion of arbitrary coupling strength $g_{\gamma,i}$ existing anywhere within the haloscope’s scan range.¹³⁸ In the common case where the data indicates the absence of an axion, the Bayesian language of updated belief maps onto the language of frequentist exclusion (Ref. [235] and Secs. 7.3.2 and 6.3.2).

The BPM framework is readily motivated by the fact that both the FT (Sec. 6.3) and BT (Sec. 7.3) frameworks discard valuable information. Each power excess, measured on a continuum, is reduced to a simple “click” or “no-click” response of an effective binary detector. A measurement that comes in just below threshold on a higher-order rescan after already exceeding threshold on previous scans is orders of magnitude more likely to indicate the presence of an axion than is a typical, 0σ power excess. Yet the threshold frameworks treat these events identically.

The BPM framework uses Bayes’ theorem to account for the precise effect of any initial measured power $x_i^{(1)}$ on the prior probability $P_{a,i}^{(0)}$ that an axion resides in bin i . After measurement, the prior is updated to the posterior,

$$P_{a,i}^{(1)} = u_i^{(1)} P_{a,i}^{(0)}, \quad (7.2)$$

where $u_i^{(1)}$ denotes the first scan’s prior update. If subsequent scans $j > 1$ are performed, $P_{a,i}^{(j-1)}$ will be further updated to

$$P_{a,i}^{(j)} = u_i^{(j)} P_{a,i}^{(j-1)}, \quad (7.3)$$

and so forth. In the appropriate limit of infinitesimal priors (Sec. 7.1.3), the single-scan prior update

¹³⁸ In the language of Bayesian statistics, a change in probability is synonymous with a change in “belief.”

for bin i and scan j follows from Eq. (7.1), independent of one's choice of prior:

$$u_i^{(j)} \approx \frac{P(x_i^{(j)}|\mathcal{A}_i)}{P(x_i^{(j)}|\mathcal{N}_i)} = \exp \left[-\frac{(\mu_{a,i}^{(j)})^2}{2} + \mu_{a,i}^{(j)} x_i^{(j)} \right]. \quad (7.4)$$

The second equality uses the Gaussian probability densities of Eq. (6.15), and \mathcal{A}_i or \mathcal{N}_i denotes the event that an axion does or does not reside in bin i , respectively. The approximation in Eq. (7.4) identifies the prior update with the Bayes factor, which compares the likelihood of two hypotheses; it applies when the axion and other sources of excess power are sufficiently unlikely.¹³⁹ The assumption underlying the validity of this approximation is made quantitative in Sec. 7.1.3, and in Sec. 7.1.5 we show that this assumption is conservative with respect to axion exclusion. The BPM analysis can therefore alternatively be viewed as a Bayes factor analysis, but the interpretation of the Bayes factor as a conservative approximation to the prior update makes for a more intuitive and immediately useful end result. Equation (7.4) also indicates that the update to the probability of an axion's existence, plotted in red in Fig. 6.1, scales exponentially with the measured power excess $x_i^{(j)}$. Rather than assume a particular axion distribution, we now consider the prior update $u_i^{(j)}$ as a function of $\mu_{a,i}^{(j)}$. Formally, $u_i^{(j)}$ is maximized at $\mu_{a,i}^{(j)} = x_i^{(j)}$. Thus at each frequency where a positive excess $x_i^{(j)} > 0$ is measured, we can single out a maximum likelihood (ML) axion distribution parametrized by $\mu_{a,i}^{(j)} = x_i^{(j)}$.¹⁴⁰ The update to the prior probability in the ML axion scales sharply with $\mu_{a,i}^{(j)} \propto g_\gamma^2$ as

$$u_{i,\text{ML}}^{(j)} = \exp \left[\frac{(\mu_{a,i}^{(j)})^2}{2} \right]. \quad (7.5)$$

Since the updates $u_i^{(j)}$ can be mapped back to their measured power excesses $x_i^{(j)}$ and sensitivities $\eta_i^{(j)}$, this use of Bayes' theorem preserves the information content of each individual measurement.

Applying Eq. (7.4) to all bins $x_i^{(j)}$ across all scans yields a spectrum of total updates U_i , defined as

$$U_i = \frac{P_{a,i}^{(M_i)}}{P_{a,i}^{(0)}} = \prod_{j=1}^{M_i} u_i^{(j)}, \quad (7.6)$$

¹³⁹ For a discussion of how Bayes factors relate to some other measures of evidence against a probabilistic hypothesis, see Ref. [245]. For discussion of the properties of related quantities in rejecting hypotheses, see Refs. [246–248].

¹⁴⁰ Conversely, measurement of a negative excess $x_i^{(j)} < 0$ reduces the probability of any axion relative to the no-axion hypothesis.

where M_i is the total number of scans performed on bin i . Unlike in the FT framework of Sec. 6.3, scans that are not performed do not enter into the equation, and the number of scans need not be specified in advance. Subject to its equivalence to the Bayes factor (Secs. 7.1.3 and 7.1.5), each total update U_i is the change in probability that an axion of a specified coupling g_γ resides in the bin. For example, $U_i = 0.1$ implies that it is 10% as probable that an axion resides in the bin after all M_i scans were performed than it was prior to the experiment. A total prior update of $U_i = 0.1$ carries exactly the same meaning as 90% exclusion ($E = 0.9$) in the standard analysis framework, in the limit of vanishing false-positive rates (Sec. 7.3.2). Therefore the equivalent 90% exclusion region consists of the (frequency-dependent) couplings above those where $U_i = 0.1$.¹⁴¹

7.1.2 Aggregation and the look-elsewhere effect

The individual $P_{a,i}^{(j)}$, or U_i , regarded as functions of g_γ , can be spectrally combined into an aggregate prior update¹⁴² function:

$$\mathcal{U}(g_\gamma) = \frac{\mathcal{P}'_a(g_\gamma)}{\mathcal{P}_a(g_\gamma)} = \frac{\sum_{i=1}^N P_{a,i}^{(M_i)}(g_\gamma)}{\sum_{i=1}^N P_{a,i}^{(0)}(g_\gamma)} \approx \frac{1}{N} \sum_{i=1}^N U_i(g_\gamma), \quad (7.7)$$

where $\mathcal{P}'_a(g_\gamma)$ and $\mathcal{P}_a(g_\gamma)$ are the aggregate posterior and prior probabilities, respectively, of that axion existing. The aggregate update $\mathcal{U}(g_\gamma)$ does for the entire haloscope run what the total update $U_i(g_\gamma)$ does for each bin: it expresses the update to the probability that an axion of coupling g_γ resides anywhere within the scanned frequency window. The final equality of Eq. (7.7) holds in the limit of approximately uniform priors,

$$P_{a,i}^{(0)} \approx \frac{\mathcal{P}_a}{N}, \quad (7.8)$$

¹⁴¹ As discussed at length in Sec. 6.5, care should be taken in interpreting this region on an exclusion plot. A region which accounts for the look-elsewhere effect with respect to exclusion can be generated above where the aggregate prior update \mathcal{U} of Sec. 7.1.2 is 0.1. However, even this region should, by the logic of the previous chapter, be aggregated with other such regions in order to properly gauge what parameter space has been excluded.

¹⁴² In this text, we refer to four kinds of prior update. The *single-scan* updates $u_i^{(j)}$ of Eq. (7.4) are multiplicatively combined into the *total* updates of Eq. (7.6) at each frequency. The total updates are subsequently aggregated via Eq. (7.7) into the *aggregate* update. The *subaggregated* updates of Fig. 7.3a are likewise aggregated, but over 1% of the HAYSTAC search window apiece. All updates are functions of axion-photon coupling g_γ , and are written explicitly as such where relevant. The quantities as calculated are Bayes factors, identifiable as updates subject to Secs. 7.1.3 and 7.1.5.

applicable when the total frequency scan range is small compared to the lowest scanned frequency, $\nu_N - \nu_1 \ll \nu_1$, as is the case for HAYSTAC or ADMX. For low-frequency experiments such as ABRACADABRA [110] or DM Radio [244], or the axion-like particle (ALP) search of Ch. 9, logarithmically uniform priors, Eq. (7.14), should be used [95].

Regardless of the measured excess powers $x_i^{(j)}$ and sensitivities $\eta_i^{(j)}$, there will exist a coupling $g_{\gamma,\text{low}}$ below which virtually nothing is learned about the presence or absence of an axion,

$$\mathcal{U}(g_\gamma \leq g_{\gamma,\text{low}}) \approx 1. \quad (7.9)$$

Likewise, there will exist another coupling $g_{\gamma,\text{high}}$ above which the probability of an axion existing vanishes

$$\mathcal{U}(g_\gamma \geq g_{\gamma,\text{high}}) \approx 0. \quad (7.10)$$

Between these two extremes, $\mathcal{U}(g_\gamma)$ will depend strongly on the measured power spectra. In the context of a realistic search, for which the experimental overhead is high, $\mathcal{U} > 1$ suggests that at least some bins within the range should be rescanned until either a cause of excess power is found or \mathcal{U} regresses to a lower value.

The aggregate prior update has several features that account for the look-elsewhere effect — the linear growth in expected number of locally significant excesses with the number of independent hypotheses tested discussed in Sec. 6.5 (for dependent hypotheses, see Sec. 7.1.4). Because the denominator of Eq. (7.7) also grows linearly with tests performed, these two linear factors cancel out. Secondly, the typical coupling excluded will move upwards (Sec. 7.3.3) as more independent tests are conducted, reflective of the increasing difficulty in ruling out a special frequency when it has more imposters to hide among (Footnote 130). Finally, the ability to incorporate nonuniform priors is a key feature for experiments sensitive over fractionally large spectral windows. Since 90% of all independent axion tests in such an experiment will occur in the highest decade (assuming equal bin widths), a standard frequentist trials factor approach [235] would set an artificially higher bar for discovery in lower decades than it would have if the highest decade were not scanned; logarithmic priors scale the height of that bar with the number of tests performed.

7.1.3 Choice of priors

For any choice of prior belief $P_{a,i}^{(0)}$ in axionic dark matter, Bayes' theorem can be applied and updates U_i calculated and reported. However, the utility of the BPM framework rests on the fact that its prior updates U_i can be treated as independent of the priors chosen for the axion and competing hypotheses, i.e. as reasonably approximated by Bayes factors. Additionally, as argued in Sec. 7.3.2, it is highly convenient that in the appropriate limit the prior updates of the BT2 framework, itself directly comparable to the BPM framework, become the conditional probabilities quoted in frequentist exclusion. These conveniences hold only so long as one's priors are infinitesimal but nonzero.¹⁴³ In this section, we will justify the assumption of infinitesimal priors, and consider a toy estimate of a prior for illustrative purposes. Additionally, we will discuss how the prior update can be used to directly inform scan strategy in the presence of nonaxionic radio frequency (RF) power excesses.

The constraint that the prior be infinitesimal is only used in simplifying the denominator of Bayes' theorem, Eq. (7.1), according to the law of total probability:

$$\begin{aligned} P(x_i^{(j)}) &= \int_0^\infty P(x_i^{(j)}|\mathcal{A}_i(g_\gamma)) P(\mathcal{A}_i(g_\gamma)) dg_\gamma + P(x_i^{(j)}|\mathcal{N}_i) P(\mathcal{N}_i) \\ &\approx P(x_i^{(j)}|\mathcal{N}_i), \end{aligned} \tag{7.11}$$

where the axion hypothesis $\mathcal{A}_i(g_\gamma)$ is here written explicitly as a family of hypotheses parametrized by $g_\gamma > 0$. The case of $g_\gamma = 0$ is indistinguishable from the no-axion hypothesis \mathcal{N}_i from the perspective of an axion haloscope.¹⁴⁴ The approximation in Eq. (7.11) is justified so long as the total probability of there being an axion in any bin i ,

$$P(\mathcal{A}_{i,\text{tot}}) = \int_0^\infty P(\mathcal{A}_i(g_\gamma)) dg_\gamma, \tag{7.12}$$

is considerably smaller than any prior update that would be applied according to Eqs. (7.4) and (7.6).

¹⁴³ The fact that the prior must be nonzero is justified by Cromwell's rule, which states that priors of exactly 0 or 1 should be avoided generally. More directly, there would be no sense in performing a haloscope search if the chance of success was truly zero.

¹⁴⁴ The possibility that dark matter more generally might not couple to electromagnetism, or any of the non-gravitational forces, is a distinct and disquieting one.

To make a coarse estimate of $P(\mathcal{A}_{i,\text{tot}})$, we can write down a Drake-like equation for axionic dark matter to which a typical haloscope searching for QCD axions such as HAYSTAC or ADMX might be sensitive:

$$P(\mathcal{A}_{i,\text{tot}}) = P(\text{PQ}) \times P(\text{DM}) \times P(m_a \approx h\nu_i/c^2), \quad (7.13)$$

where PQ denotes the event of the Peccei-Quinn hypothesis (Sec. 2.3.2) being correct and the QCD axion field existing, DM denotes the event of that axion field actually accounting for an appreciable fraction of the galaxy's dark matter (Sec. 2.4.2), $m_a \approx h\nu_i/c^2$ denotes the event that the axion is within roughly a linewidth of the bin of interest, i , and the probability of each event in parentheses is implicitly conditional upon all leftward events. In practice several other propositions must be true for an axion to be detectable. For instance, the axion must not be coincident with any other RF spikes and its lineshape should roughly match that used in the data processing (Sec. 2.4.3).

It is beyond the scope of this thesis to attempt serious estimates of these probabilities, but for the sake of justifying the infinitesimal prior approximation, we only need to consider the sample optimistic set of probabilities shown in Table 7.1. For logarithmically uniform distributed prior belief [95]

$$P_{a,i} \propto \frac{1}{\nu_i} \quad (7.14)$$

in a $Q_a \sim 10^6$ axion between $m_{\text{SN}} \sim 10 \mu\text{eV}$ and $m_{\text{OC}} \sim 1 \text{ meV}$, the approximate mass range not disfavored by evidence from the SN1987A neutrino burst or overclosure arguments (Sec. 2.4.1) [52–54],

$$P(m_a \approx h\nu_i/c^2) = \frac{\log(1 + Q_a)}{\log(m_{\text{SN}}/m_{\text{OC}})} = 2.1 \times 10^{-7}. \quad (7.15)$$

If overclosure arguments do not apply [249], then a QCD axion could be far lighter than m_{OC} . Moreover, in practice, the prior probability for an axion to be in any given narrow coupling window must be less than $P(\mathcal{A}_{i,\text{tot}})$. Hence, the rough estimate of

$$P(\mathcal{A}_{i,\text{tot}}) \sim 9 \times 10^{-8} \quad (7.16)$$

should be considered a generous upper bound on the prior probability as the term is used in this thesis.

event	optimistic probability
PQ	0.8
DM	0.5
$m_a \approx h\nu_i/c^2$	2.1×10^{-7}
$\mathcal{A}_{i,\text{tot}}$	9×10^{-8}

Table 7.1: Sample set of probabilities used to estimate prior axion hypothesis probability according to Eq. (7.13). The probabilities $P(\text{PQ})$ and $P(\text{DM})$ are set to high (optimistic) values, where it is clear that raising them would not qualitatively change the conclusion that the priors are in the appropriate infinitesimal limit. The probability of the axion mass coinciding with a given bin assumes logarithmic priors [95], and is itself perhaps optimistic in assuming that the axion exists within a three-mass decade window in which the haloscope operates.

In practice, one should not wait for a prior update of order $1/P(\mathcal{A}_{i,\text{tot}})$ to manually interrogate a bin;¹⁴⁵ instead, an approximate prior estimate of the prevalence of spurious RF tones (Sec. 7.1.5) should inform decision making. The prevalence of these RF spikes is *a priori* unknown, but the fact that in practice [78, 105, 107] the number of rescans required agrees with the predictions from Gaussian statistics (i.e. the measured distribution of excess powers looks Gaussian even in its right tail) indicates that only a few RF spikes large enough to push a bin above threshold are expected in a scan. For an initial scan with N_I independent bins, therefore, a prior update of $U \sim N_I/N_{\text{RF}}$ should be sufficient cause to manually interrogate the bin. Here, N_{RF} is the experimentalist’s best guess for the expected number of real RF spikes in the scan window. Insofar as this guess is incorrect, the haloscope search will take longer than would be optimal. Rescanning and eliminating spurious RF tones ensures that the bias in the updates remains small, as well as conservative.¹⁴⁶

The prior probability $P(\mathcal{A}_{i,\text{tot}})$ estimated in this section makes our approximation of infinitesimal priors valid for any single-bin updates $U \ll 1/P(\mathcal{A}_{i,\text{tot}})$. Since the prior update at which manual interrogation ought to begin, N_I/N_{RF} , is generally orders of magnitude less than $1/P(\mathcal{A}_{i,\text{tot}})$, the approximation of infinitesimal priors will always be valid.

¹⁴⁵ Manual interrogation means performing tests that would discriminate between most spurious RF excesses and an axion: for example tuning the spike well off cavity resonance and/or ramping down the magnet and seeing if the signal persists. See Item 9 in Sec. 6.2.1.

¹⁴⁶ The conservative bias discussed in Sec. 7.1.5 results from interpreting the no-axion hypothesis, \mathcal{N}_i in Eq. (7.11), to incorporate spurious RF tones. This makes $P(x_i^{(j)}|\mathcal{N}_i)$ larger in reality than as calculated from Eq. (6.15). Using a slightly undersized denominator in Eq. (7.4) results in a slight anti-exclusion bias.

7.1.4 Correlations between grand spectrum bins

In this section, we will see that correlations between grand spectrum powers¹⁴⁷ at different frequencies do not bias the BPM framework’s outputs, and are furthermore desirable. The majority of Sec. 7.1 treats the grand spectrum power excesses as uncorrelated for simplicity. In practice, this assumption is violated, as grand spectrum bin spacings are intentionally made spectrally smaller than the axion linewidth and the frequency window of the digital Savitzky-Golay (SG) filter applied to the data. The effect of estimating power excesses based on the axion lineshape (Step 8 in Sec. 6.1) is to positively correlate nearby bins, whereas the effect of the SG filtering (Step 4) is to negatively correlate bins over a somewhat longer frequency scale [78]. These are the only two processing steps that will correlate adjacent bins.¹⁴⁸ The data processing introduces no interscan correlations, as it handles separate scans independently. The single-scan prior updates $u_i^{(j)}(g_\gamma)$ can therefore be treated as unconditionally unbiased ($\langle u_i^{(j)} \rangle = 1$), scan-independent, frequency-dependent random variables.

Given these properties of the single-scan prior updates $u_i^{(j)}(g_\gamma)$, the result that the aggregate prior update $\mathcal{U}(g_\gamma)$ is unbiased then follows from the fact that aggregation entails multiplying across scans and summing across bins. The total prior updates $U_i(g_\gamma)$ are defined in Eq. (7.6) as the product of grand spectrum prior updates from M_i scans. Using the fact that the expectation value of the product of *independent* random variables is the product of the expectation values, here

$$\langle U_i(g_\gamma) \rangle = \left\langle \prod_{j=1}^{M_i} u_i^{(j)}(g_\gamma) \right\rangle = \prod_{j=1}^{M_i} \langle u_i^{(j)}(g_\gamma) \rangle = 1, \quad (7.17)$$

the total prior updates are seen to be unbiased, i.e. have unit expectation. Aggregation then occurs through summation of the posteriors and priors via Eq. (7.7). The expectation value of the sum of

¹⁴⁷ Recall from Sec. 6.1 that the grand spectrum is the final spectrum resulting from the data processing, accounting for noise level, signal strength, and axion lineshape.

¹⁴⁸ Systematic effects in the measurement will also correlate grand spectrum power excesses. The identification and removal of nonaxionic power excesses and the SG filtering exist largely to undo such experimentally-induced correlations. We do not consider the effect of such correlations on the grounds that these processing steps appear largely successful [78].

dependent random variables is the sum of the expectation values, here

$$\langle \mathcal{U}(g_\gamma) \rangle = \left\langle \sum_{i=1}^N U_i(g_\gamma) \right\rangle = \sum_{i=1}^N \langle U_i(g_\gamma) \rangle = 1, \quad (7.18)$$

where the last equality uses the result of Eq. (7.17) to demonstrate that $\mathcal{U}(g_\gamma)$ is also unbiased. As we relax assumptions of uniform or even low priors, the relevant updates, treated as random variables, can be shown to have unit expectation independent of grand spectrum correlations, but the formulae used to compute and combine the updates become less elegant.

The primary effect of grand spectrum correlations on the prior updates is not to bias the analysis, but only to alter the higher moments of the distribution that aggregate prior updates are drawn from. In particular, the number N_I of *independent* grand spectrum bins is smaller than the number N of grand spectrum excesses calculated. This increases the spread of the distribution that aggregate prior updates will be drawn from for a given haloscope scan window. However, since the dominant correlations are produced by the axion lineshape (Step 8, Sec. 6.1), the number of independent bins has as a rough lower bound the number of axion linewidths Δ_a in the scan window Δ_W , and as an upper bound the number of bins used: $\Delta_W/\Delta_a \lesssim N_I \leq N$. The fact that $N_I \leq N$ broadens the probability distributions of the aggregate outcomes (Sec. 7.3.3), but does not otherwise impact the analysis.

To see that correlations are indeed desirable, note that their absence would imply a grand spectrum bin spacing larger than the axion linewidth. As discussed in Ref. [78], such a coarse bin spacing would effectively leave unprobed a large fraction of the scan window. Since the correlations produce no deleterious effects upon the BPM framework, there is no reason not to oversample the space of possible axion masses, and the correlations should be considered a feature, not a bug.

7.1.5 Practical departures from Gaussianity

The formula for updating priors using the BPM framework, Eq. (7.4), assumes that the axion and no-axion distributions are Gaussian, Eq. (6.15). In particular, measured distributions often rise in their extreme tails over what is ideally expected [250]. In the case of haloscopic axion

detection, the presence of spurious RF tones increases the probability density of the high-power tail of the probability density functions (PDFs) of Fig. 6.1. In this section, we will see that the interpretation of the exclusion achieved with the BPM framework is robust to the resultant bias towards discovery. Conversely, there is little to be gained from Bayesian inference in the context of manually interrogating a persistent signal with the goal of reporting discovery. Efficiency is a far less important consideration in this context, and test statistics such as those of Refs. [75, 241], in conjunction with manual interrogation (Item 9 of Sec. 6.2.1 and Footnote 145) are well suited to the task of claiming discovery, once all manual interrogation tests are passed.

Sources of RF noise in the haloscope’s physical environment can couple into the receiver chain. Real haloscopes go to significant effort to discriminate such artificial excesses from real axion signals, but excesses sufficiently similar in spectral profile to an axion necessarily contribute to the measured power distributions, leading to an effective increase in the high power probability density of both the axion and no-axion distribution PDFs.

The effect of increased probability density in the high-power tail in either distribution in Fig. 6.1 is to bias the BPM analysis towards discovery of an axion. In any one frequency bin, this effect is very unlikely because it only affects the distribution tails. Indeed, for existing haloscopes, the measured power excesses seem very well approximated out to several standard deviations by Gaussian distributions [78, 125], indicating that the consequences of non-Gaussian features will be modest. However, with tens or hundreds of thousands of independent frequency bins measured, there may be a small augmentation of the BPM framework’s aggregate prior update, Eq. (7.7).

In the context of reporting exclusion over axion parameter space — the only context in which data from haloscope searches has appeared to date — a bias towards discovery is conservative, and hence acceptable. In the context of an FT framework (Sec. 6.3), the effect of a discovery bias is not to change the reported exclusion, but rather to increase expected rescan time and the chance of a false positive. Since, as we will see in Fig. 7.3b, the BPM framework still achieves a 36% scan rate enhancement relative to thresholding for the HAYSTAC Phase I dataset, the conservative bias is manifestly not large enough (at least for this dataset, and also the Phase II dataset of Ch. 8) to

overwhelm the improved exclusion over the whole scan range.

In the context of reporting the discovery of axionic dark matter, conversely, a bias towards discovery would appear to be harmful. However, the nature of axion detection guards against adverse outcomes in practice. A power excess presenting as an axion, once identified, is fairly straightforwardly shown to persist or not by manual interrogation. The role of the statistical analysis frameworks discussed in this thesis is not to firmly identify the axion as such once a persistent signal is found, but rather to indicate with the highest fidelity where such a signal is likely to lie. While a BPM analysis may bias the experimentalist towards believing that a few nonaxionic excesses might be axionic, the real certification of an axion signal would look like that for other discoveries in fundamental physics: a series of careful measurements precluding beyond doubt that no other known physical phenomena could account for the signal. In other words, the asymmetry between the criteria for exclusion and discovery ensure that BPM’s discovery-bias would almost never¹⁴⁹ lead to accidental discovery.

7.2 Effects of data processing on information content

This chapter takes as its motivation the unacceptably long scan times that will likely be required to verify or falsify the axionic dark matter hypothesis over any meaningful fraction of the plausible parameter space. However, the BPM framework optimizes only the statistical analysis of the already-processed data of a specific measurement, raising two complementary questions: is the measurement optimal, and is the data processing optimal? The first question, of whether the measurement itself is optimal — i.e. is the haloscope as typically operated the most efficient tool for evaluating the axionic dark matter hypothesis? — is beyond the scope of this thesis, and the reader is referred to Ref. [95]. Subject to a haloscope or haloscope-like search platform, however, this section addresses the effect of the data processing — the intermediate set of steps

¹⁴⁹ It is possible to construct scenarios where nature is sufficiently cruel as to source a tone with the distinguishing properties of the axion. For example, imagine a microwave tone generated by the magnetic field itself which manages to vibrate a piezoelectric object attached to the outside of the cavity, which in turn creates eddy currents inside it. The result is a B_0^2 -proportional microwave tone inside the cavity. Even still, however, it is unlikely (thought not impossible) that such a tone would mimic the expected axion lineshape, Eq. (2.14).

between measurement and statistical analysis detailed in Sec. 6.1 — on the information content of the haloscope data. We will see that the current data processing protocol of the leading haloscope experiments [91, 107] are already highly optimized, with very little room for further improvements, though this has not historically been the case.

This section evaluates the consequences of each data processing step enumerated in Sec. 6.1 on the axion-pertinent (AP) information content of the data, where AP is understood to mean “potentially bearing on the probability of the axion hypothesis being true.” Information pertinent to the existence of the axion can only be generated during the measurement, and may be degraded by the subsequent data processing. Our aim is to determine where in Steps 1–10, if at all, AP information could be lost. A guaranteed way to make use of all of the AP information content would therefore be to express the predictions of the axion and no-axion hypotheses, $\mathcal{A}_{Q,ijk}$ and $\mathcal{N}_{Q,ijk}$, respectively, in the space of the quadrature voltages $v_{Q,ijk}^\alpha$ and apply Bayes’ theorem as

$$u_{Q,ijk}^{\alpha,\text{Re}} = \frac{P(\text{Re}[v_{Q,ijk}^\alpha]|\mathcal{A}_{Q,ijk})}{P(\text{Re}[v_{Q,ijk}^\alpha]|\mathcal{N}_{Q,ijk})}, \quad (7.19)$$

where the standard infinitesimal prior limit has been assumed and the same update applies for the imaginary components $\text{Im}[v_{Q,ijk}^\alpha]$.¹⁵⁰ While this approach is impractical, it nonetheless serves as an excellent baseline to compare the effects of the data processing protocol against.

Step 1 includes two operations: squaring the real and imaginary parts of each quadrature voltage, and then adding them. The real and imaginary parts of each quadrature voltage can be approximated as the same independent, identically distributed Gaussian random variables with mean zero [75, 200]. Absent an axion, we label the variances of these random variables $\sigma_{v,ij}^2$. An axion, assumed to be of a given coupling and at a given frequency, increases the variance to $\lambda_{ij}\sigma_{v,ij}^2$, where λ_{ij} is extremely close to but greater than unity. The prior update, Eq. (7.19), is the ratio of Gaussian PDFs given by Eq. (6.15):

$$u_{Q,ijk}^{\alpha,\text{Re}} = \frac{1}{\sqrt{\lambda_{ij}}} \exp \left[\frac{1 - \lambda_{ij}^{-1}}{2\sigma_{v,ij}^2} (\text{Re}[v_{Q,ijk}^\alpha])^2 \right]. \quad (7.20)$$

¹⁵⁰ While in principle it is possible to evaluate the axionic dark matter hypothesis in the time domain, we do not expect that this would provide any benefit, given the reversibility of the Fourier transform and the fact that the phase of the axion field at any moment is unknown.

Since $\text{Re}[v_{Q,ijk}]$ only occurs squared in this prior update (likewise $\text{Im}[v_{Q,ijk}]$ in its identical equation), no AP information is lost by squaring the real and imaginary quadrature voltage components. Furthermore, since Bayesian updates are by their nature combined multiplicatively, the update accounting for both components of the quadrature voltage is

$$u_{Q,ijk}^{\alpha,v} = \frac{1}{\lambda_{ij}} \exp \left[\frac{1 - \lambda_{ij}^{-1}}{2\sigma_{v,ij}^2} \left((\text{Re}[v_{Q,ijk}^\alpha])^2 + (\text{Im}[v_{Q,ijk}^\alpha])^2 \right) \right]. \quad (7.21)$$

Because the voltage quadrature components appear in the prior update only in their Eq. (6.1) combination, Step 1 preserves AP information content.

The same argument applies to Steps 2 and 3. Multiplicative prior updates have their powers add in the exponent, hence Steps 2 and 3 do not degrade AP information content. For Step 2, this argument assumes that the axion and the thermal noise fluctuations distribute power evenly between the X and Y quadratures. For a two-quadrature measurement where the phase of the axion is unknown, that assumption is valid.

Step 4 is among the more vulnerable steps to meaningful information loss. If $\langle x_{ij}^\beta \rangle$ in Eq. (6.4) is treated as a number, then Eq. (6.4) amounts merely to the division and subtraction of constants from the raw spectrum random variables x_{ij}^β , which will not degrade the information content. In practice, however, $\langle x_{ij}^\beta \rangle$ must itself be estimated, making it in effect a random variable. The estimation used by HAYSTAC and ADMX is performed via SG filters, which estimate $\langle x_{ij}^\beta \rangle$ at each frequency bin i via the polynomial generalization of a moving average. If a wide spectral window for the generalized moving average is used, then $\langle x_{ij}^\beta \rangle$ is estimated with low variance, and acts like a constant. However, the SG filtering creates undesired correlations between bins up to two window-lengths apart, while also slightly attenuating an axion's visibility, an effect accounted for in Step 9. The trade-off between the desired filtering effects and the undesired correlations and attenuation is beyond the scope of this thesis, but is discussed at length in Ref. [78]. So far, no proof exists that SG is the optimal choice of digital filter¹⁵¹ for the data, so Step 4 may admit of

¹⁵¹ Recall that the SG is itself a low pass filter, which the spectra are divided by so as to preserve high frequency structure. See Ref. [78] for numerical tests performed to quantify the effects on synthetic data with real haloscope spectral structure.

meaningful optimization.

Calculating the processed spectrum sensitivity parameters η_{ij}^γ of Step 5 relies on independent measurements of the thermal and added noise, magnetic field, and cavity quality factor and mode-structure properties of the haloscope, as well as calculations of the power that a dark matter axion would deliver. Subject to the accuracy of these measurements and calculations, Step 5 simply divides the processed spectra power excess random variables x_{ij}^γ by constants, preserving information content.

The ML-weighting of Step 6 is the information-preserving generalization of the straightforward addition of Steps 1–3, for the case where the random variables being added have different variances. By the central limit theorem, the rescaled power excesses x_{ij}^δ are Gaussian-distributed with known variances σ_{ij}^δ well approximated as independent of the presence of any axion which delivers power far less than vacuum. The pertinent effect of an axion is to shift the mean of the rescaled power excess from 0 to μ^δ , which has no bin- or tuning step-dependence by construction. The appropriate prior update delivered by each rescaled power excess x_{ij}^δ is thus

$$u_{ij}^\delta = \exp \left[\frac{-(\mu^\delta)^2/2 + \mu^\delta x_{ij}^\delta}{(\sigma_{ij}^\delta)^2} \right]. \quad (7.22)$$

When the updates of all N_t tuning steps are multiplied, the result is

$$\begin{aligned} \prod_{j=1}^{N_t} u_{ij}^\delta &= \exp \left[\sum_{j=1}^{N_t} \left(\frac{-(\mu^\delta)^2/2 + \mu^\delta x_{ij}^\delta}{(\sigma_{ij}^\delta)^2} \right) \right] \\ &= \exp \left[\frac{-(\mu^\delta)^2/2 + \mu^\delta x_i^\epsilon}{(\sigma_i^\epsilon)^2} \right]. \end{aligned} \quad (7.23)$$

Because the updates from all bin- i rescaled spectra power excesses x_{ij}^δ and variances $(\sigma_{ij}^\delta)^2$ can thus be obtained using the single bin- i combined spectrum power excess x_i^ϵ and variance $(\sigma_i^\epsilon)^2$, given by Eqs. (6.6) and (6.7), respectively, the ML estimation of Step 6 preserves AP information content.

The same argument about ML estimation guarantees that Steps 7 and 8 preserve AP information as well, in the limit where the axion lineshape is truly approximately constant over spectral scales of $n_c \Delta_b$. Step 9 is in effect properly accounting for the information loss suffered during Step 4, but is itself simply multiplication by a scalar, and causes no additional information loss.

7.3 Comparison of analysis frameworks

One of the central claims of this chapter is that in a haloscope analysis an apples-to-apples comparison may be made between the widely quoted frequentist exclusion E of Sec. 6.3 and the Bayesian prior update U_i of Sec. 7.1.1. In general, these two are in fact not equivalent. The equivalence holds only between FT analyses with low false-positive rates and Bayesian analyses using low priors. In this section, I introduce two Bayesian threshold analyses. The first framework, BT1 (Sec. 7.3.1), is more informative than the second, BT2 (Sec. 7.3.2). In the dual limit of infinitesimal priors/low false-positive rates, BT2 reveals the Bayesian prior update to be completely equivalent to the frequentist exclusion. In Sec. 7.3.3, we will quantitatively examine the outcomes of all three Bayesian frameworks (BT1, BT2, and BPM) discussed in this thesis. Together, BT1 and BT2 motivate the BPM framework as the maximally informative, minimally constrained limit of haloscope analyses.

7.3.1 Bayesian threshold 1 framework

The first of the threshold-based Bayesian analyses we consider, BT1, updates its priors after each successive scan. We are interested in the conditional probability that there is an axion in the i^{th} bin, given that the measured power in the j^{th} scan either did ($x_i^{(j)} \geq x_T^{(j)}$) or did not ($x_i^{(j)} < x_T^{(j)}$) exceed threshold. For simplicity of notation, we will drop all frequency indices i for the remainder of this and other sections where not necessary, with it understood that results can ultimately be aggregated via Eq. (7.7). We denote the event of a measured power exceeding (not exceeding) threshold as a binary detector going “click” (“no-click”), and we denote the event that the measured power excess came from the axion distribution (no-axion distribution) as \mathcal{A} (\mathcal{N}). The posterior probability of an axion given a click is

$$P(\mathcal{A}|\text{click}) = \frac{P(\text{click}|\mathcal{A})P(\mathcal{A})}{P(\text{click}|\mathcal{A})P(\mathcal{A}) + P(\text{click}|\mathcal{N})P(\mathcal{N})}, \quad (7.24)$$

where the denominator is equivalent to $P(\text{click})$. Identifying $P(\text{click}|\mathcal{A})$ as the single-scan true-positive rate $1 - f_n^{(j)}$ (Fig. 6.1, light blue) and $P(\text{click}|\mathcal{N})$ as the single-scan false-positive rate $f_p^{(j)}$

(Fig. 6.1, dark green), abbreviating the prior (posterior) probability that there is an axion in the bin as $P_a^{(0)} = P(\mathcal{A})$ [$P_{a,\text{cl}}^{(1)} = P(\mathcal{A}|\text{click})$], and noting that $P(\mathcal{N}) = 1 - P_a^{(0)}$, Eq. (7.24) simplifies to

$$u_{\text{cl}}^{(1)} = \frac{P_{a,\text{cl}}^{(1)}}{P_a^{(0)}} \approx \frac{1 - f_n^{(1)}}{f_p^{(1)}} \approx \frac{1}{f_p^{(1)}}. \quad (7.25)$$

The first approximation holds in the appropriate limit of infinitesimal priors, here $P_a^{(0)} \ll f_p^{(1)}$, and the second holds in the limit of low single-scan false-negative rates, $f_n^{(1)} \ll 1$.

In the event of a no-click, the posterior probability $P_{a,\text{no}}^{(1)} = P(\mathcal{A}|\text{no-click})$ is given by an expression analogous to Eq. (7.24) containing the true- ($1 - f_p^{(j)}$; Fig. 6.1, light green) and false- ($f_n^{(j)}$; Fig. 6.1, dark blue) negative rates. The expression simplifies as

$$u_{\text{no}}^{(1)} = \frac{P_{a,\text{no}}^{(1)}}{P_a^{(0)}} \approx \frac{f_n^{(1)}}{1 - f_p^{(1)}} \approx f_n^{(1)}, \quad (7.26)$$

where the first approximation holds for low priors $P_a^{(0)} \ll 1$, and the second holds for low single-scan false-positive rates, $f_p^{(1)} \ll 1$. Whereas single-scan false-negative rates in haloscope searches using an FT framework are typically close to 5%, false-positive rates are often kept much lower, to minimize the need for time-expensive rescans.

Applied to a sequence of M scans, the BT1 analysis provides a crucial operational freedom, shared by the BPM framework, to the experimentalist. Since the ordering of multiple updates is inconsequential, a series of M identical ($f_n^{(j)} = f_{n0}$, $f_p^{(j)} = f_{p0}$) scans has $M + 1$ possible outcomes, corresponding to observing $c \in \{0, 1, \dots, M\}$ clicks and $M - c$ no-clicks. Using the simplest forms of Eqs. (7.25) and (7.26), the final prior update, obtained from Eq. (7.6), is

$$U_c^{\text{BT1}} \approx \frac{(f_{n0})^{M-c}}{(f_{p0})^c}. \quad (7.27)$$

The prior is upgraded by $1/f_{p0}$ for every positive result, and downgraded by f_{n0} for every negative result.

In the BT1 framework, the threshold and false-negative rate for each scan must be set in advance, but the number of scans M need not be. This operational freedom is a consequence of the fact that the expected value of the prior update before performing a given scan is unity, and so

the decision to perform another scan or not can be made on the fly, without biasing the outcome. In contrast, in the FT framework of Sec. 6.3, adding another scan midway through a multiscan protocol and recalculating the false-negative rate for the entire protocol, the experimentalist can manipulate the expected value of the exclusion. This freedom has nothing to do with the use of frequentist versus Bayesian inference. We will see that Bayesian frameworks can be equally restrictive.

		<u>SCAN 2</u>	
		CLICK	NO CLICK
<u>SCAN 1</u>	CLICK	1) FAIL TO REJECT NULL $P(1 \mathcal{A}) = (1 - f_{n0})^2$ $P(1 \mathcal{N}) = f_{p0}^2$	2) REJECT NULL $P(2 \mathcal{A}) = (1 - f_{n0})f_{n0}$ $P(2 \mathcal{N}) = f_{p0}(1 - f_{p0})$
	NO CLICK	3) REJECT NULL $P(3 \mathcal{A}) = f_{n0}(1 - f_{n0})$ $P(3 \mathcal{N}) = (1 - f_{p0})f_{p0}$	4) REJECT NULL $P(4 \mathcal{A}) = f_{n0}^2$ $P(4 \mathcal{N}) = (1 - f_{p0})^2$

Figure 7.1: The four possible outcomes of an $M = 2$ -identical scan threshold protocol. The conditional probabilities of each outcome (1)–(4) subject to the axion (\mathcal{A}) and no-axion (\mathcal{N}) hypotheses are given in terms of true- and false-positive and -negative rates (Fig. 6.1). In both the frequentist threshold (FT) framework of Sec. 6.3 and the Bayesian threshold 2 (BT2) framework of Sec. 7.3.2, outcomes (2)–(4) (gray) are all treated identically as the negative result: in the FT (BT2) framework, the null is rejected (the prior is reduced). In this case, the exclusion E , Eq. (6.17), is trivially equivalent to the prior update $U_{\text{neg}}^{(\text{BT}2)} \approx F_n$ in the dual limit of low false-positive rates and priors. Only region (1) (red) fails to reject the null (FT), or increases the prior (BT2). Figure reproduced from Ref. [126].

7.3.2 Bayesian threshold 2 framework

Whereas the BT1 framework offers $M + 1$ possible outcomes for M scans and provides interscan operational freedom to the experimentalist, the second BT framework (BT2) that we consider is designed to be uninformative and restrictive by comparison; its outcomes, tellingly, will

map precisely onto those of the FT framework of Sec. 6.3. The standard FT analysis permits two possible results at each frequency, illustrated for a simple case in Fig. 7.1: a rejection of the null (negative result) or a failure to reject it (positive result). For an M -scan protocol, the negative result is defined as the case where at least one scan no-clicks; the positive result occurs only when all M scans click. The multiscan false-negative rate, F_n from Eq. (6.20), sets the exclusion via Eq. (6.17). False-positive rates do not enter directly into the FT framework’s reported exclusion.¹⁵²

The BT2 framework applies Bayes’ theorem, Eq. (7.1), to this binary result landscape, obtaining a posterior either for a positive result of $P_{a,\text{pos}}^{(1)}$ or for a negative result of $P_{a,\text{neg}}^{(1)}$. For the positive case,

$$P(\mathcal{A}|\text{pos}) = \frac{P(\text{pos}|\mathcal{A})P(\mathcal{A})}{P(\text{pos}|\mathcal{A})P(\mathcal{A}) + P(\text{pos}|\mathcal{N})P(\mathcal{N})} \quad (7.28)$$

yields

$$U_{\text{pos}}^{\text{BT2}} = \frac{P_{a,\text{pos}}^{(1)}}{P_a^{(0)}} \approx \frac{1 - F_n}{F_p} \approx \frac{1}{F_p}, \quad (7.29)$$

where the total false-negative rate F_n is that of Eq. (6.20) and

$$F_p = \prod_{j=1}^M f_p^{(j)} \quad (7.30)$$

is the total false-positive rate for the M scans. The first (second) approximation in Eq. (7.29) is valid for infinitesimal priors, $P_a^{(0)} \ll F_p$ (low total false-negative rate, $F_n \ll 1$). Note that BT2’s positive prior update agrees in the case of identical scans (the general cases agree as well) precisely with the $c = M$ case of Eq. (7.27) for BT1. The agreement is due to the fact that they derive from the same event — a series of M clicks. Since false-positive rates are very small, $U_{\text{pos}}^{\text{BT2}}$ is a very large number, corresponding to just how unlikely a positive result would be to observe, without an axion present.

The negative result yields an update that will agree precisely with the standard FT exclusion

¹⁵² More completely, false-positive rates only enter indirectly, as the global false-positive rate — the chance of at least one bin achieving a positive result (i.e. a click on all M scans) — must be set sufficiently low to account for the look-elsewhere effect, as discussed in Sec. 6.5.1. Choosing a higher false-positive rate will achieve a deeper exclusion if no positive results are recorded, but will increase the probability of such an event.

E. From the equation for $P(\mathcal{A}|\text{neg})$ analogous to Eq. (7.28), we obtain:

$$U_{\text{neg}}^{\text{BT2}} = \frac{P_{a,\text{neg}}^{(1)}}{P_a^{(0)}} \approx \frac{F_n}{1 - F_p} \approx F_n, \quad (7.31)$$

where the two approximations assume low priors $P_a^{(0)} \ll 1$ and low M -scan false-positive rates $F_p \ll 1$. With these conditions met, the BT2 negative prior update is the complement of the frequentist exclusion, $U_{\text{neg}}^{\text{BT2}} \approx 1 - E$. In the limit relevant for haloscope searches, the frequentist exclusion is therefore equivalent to the Bayesian prior update, an equivalence used in the context of other searches for new physics [235]. Whether this update is obtained through one of the BT frameworks or the BPM framework is immaterial.

Finally, a pair of differences between the two BT frameworks discussed in this section hints at two important operational advantages of the BPM framework. The BT1 framework, by allowing $M + 1$ possible results with differing outcomes (or up to 2^M for scans with nonidentical false-negative and -positive rates), has an informational advantage over the BT2 framework. The BT2 framework deliberately blinds the experimentalist to which scan(s) no-clicked, combining M of the $M + 1$ possible outcomes into a single negative result. For an experimentalist intent on discovering or excluding the axion, there can be no benefit to discarding this information. Taking the next logical step, the power-measured information from each scan need not be discarded. The BPM framework, unlike the BT frameworks, uses this information.

	<i>p</i>-value	BT2/FT	BT1	BPM
possible outcomes per bin	∞	2	2^M	∞
interscan freedom	yes	no	yes	yes
intrascan freedom	yes	no	no	yes

Table 7.2: Comparison of the Bayesian threshold frameworks (BT1 and BT2) of Sec. 7.3 with the Bayesian power-measured (BPM) framework of Sec. 7.1 and the p -value framework of Sec. 6.4 for a protocol with M possibly nonidentical scans performed. More informative frameworks map the continuum of possible measurements onto a larger number of reported outcomes. The more informative frameworks also permit the experimentalist greater freedom to alter the scan protocol without biasing the outcome. Interscan changes to upcoming scans are allowed within all but the BT2 framework (equivalently, the FT framework of Sec. 6.3), while the BPM framework even permits intrascan adjustments as information compiles. The p -value framework has the same properties as the BPM framework, but comes with a number of downsides discussed in Sec. 6.4.

The BT1 framework also proves operationally superior to BT2. Whereas the BT2 framework locks in a commitment to M scans, if necessary, BT1 allows the experimentalist to reinsert himself into the decision-making *between* scans without biasing the outcome. This interscan freedom is improved to intrascan freedom for BPM, as indicated in Table 7.2, which summarizes the informativeness and operational constraints imposed by each framework. Under BPM, the expected prior update for the next iota of power measured always being unity protects the experimentalist from inserting bias. To the degree that the data processing allows it, probability can be tracked in real time and used to inform scan-protocol decisions on the fly. This capability provides an opportunity for optimizing a BPM haloscope search algorithm, a promising direction for future analysis beyond the scope of this thesis.

7.3.3 Outcome distributions for Bayesian analyses

The reanalysis of the HAYSTAC Phase I dataset that we will consider in Sec. 7.4 indicates that the BPM framework achieves superior exclusion to thresholding in the case of one real dataset (with a second added in Ch. 8). In this section, we demonstrate that this result is typical by treating the prior updates in the BPM framework and the BT1 and BT2 frameworks discussed in Sec. 7.3 as random variables. We make an apples-to-apples comparison between the aggregate prior update probability distributions obtained from an $M = 2$ -identical scan protocol in these three frameworks. In particular, while the mean aggregate prior update $\langle \mathcal{U} \rangle = 1$ for any unbiased analysis, we will see that the *median* aggregate prior update in the BPM framework is a factor of two smaller than in either threshold framework. From this reduction in the median aggregate prior update we predict a typical scan rate enhancement of 30% for BPM relative to thresholding, consistent with the enhancement observed in the HAYSTAC Phase I dataset. En route to this final result, we obtain analytic expressions for the total prior update probability distributions in each framework. Comparing the three frameworks at this level likewise elucidates the difference between the BPM and threshold exclusion lines plotted in Fig. 7.3a.

Throughout this section, we assume an ideal, axionless haloscope dataset: the grand spectrum

excesses are independent, identically distributed random variables $x_i^{(j)} \sim N(0, 1)$. We compare prior update probability distributions for an $M = 2$ -identical scan protocol in which an initial scan is performed, followed by a rescan in each bin whose measured power excess exceeds a predetermined threshold.¹⁵³ For definiteness, we assume scan parameters similar to those used in HAYSTAC Phase I. For the BT1 and BT2 frameworks, we will first derive analytic expressions for the prior updates in a more general M -identical scan protocol, and then specify to $M = 2$.

Starting with the least informative BT2 framework, the total prior update in a single bin has probability mass function (PMF):

$$\begin{aligned} P(U^{\text{BT2}} = U_{\text{pos}}^{\text{BT2}}) &= (f_{p0})^M \\ P(U^{\text{BT2}} = U_{\text{neg}}^{\text{BT2}}) &= 1 - (f_{p0})^M. \end{aligned} \quad (7.32)$$

In any scan of N independent frequency bins, $k = 0, \dots, N$ of those will realize a positive result, where k is binomial distributed $k \sim B(N, (f_{p0})^M)$. The PMF for the aggregate prior update is then

$$P\left(\mathcal{U}^{\text{BT2}} = \frac{U_{\text{pos}}^{\text{BT2}} - U_{\text{neg}}^{\text{BT2}}}{N}k + U_{\text{neg}}^{\text{BT2}}\right) = \binom{N}{k} [(f_{p0})^M]^k [1 - (f_{p0})^M]^{N-k}. \quad (7.33)$$

Next we consider the BT1 framework, whose total prior update U_c^{BT1} , Eq. (7.25), is written in terms of the single-scan updates u_{cl} and u_{no} for a click, Eq. (7.25), and for a no-click, Eq. (7.26), respectively. Each click occurs with probability f_{p0} , hence the probability of obtaining exactly $c = \{0, 1, \dots, M\}$ clicks is

$$P_c = (f_{p0})^c (1 - f_{p0})^{1 - \delta_{Mc}}, \quad (7.34)$$

where δ_{ab} is the Kronecker delta function. The total prior update PDF is therefore

$$P\left(U^{\text{BT1}} = [u_{\text{cl}}]^c [u_{\text{no}}]^{1 - \delta_{Mc}}\right) = P_c. \quad (7.35)$$

For the M -identical scan protocol over N independent bins, the aggregate prior update, Eq. (7.7), is parametrized by the numbers n_c of bins that click c times, which are multinomial distributed

¹⁵³ As discussed in Sec. 7.3, BPM provides freedoms to deviate from such a rigid protocol, and can therefore outperform the projections here. In Ch. 8, we take advantage of the enhanced operational freedoms of the BPM framework.

with probabilities P_c given by Eq. (7.34). The aggregate prior update is

$$P\left(\mathcal{U}^{\text{BT1}} = \frac{1}{N} \sum_{c=0}^M n_c \frac{(u_{\text{cl}})^c}{(u_{\text{no}})^{\delta_{M_c-1}}}\right) = \frac{N!}{\prod_{c'=0}^M n_{c'}!} \prod_{c'=0}^M (P_{c'})^{n_{c'}} \quad (7.36)$$

if

$$\sum_{c=0}^M n_c = N, \quad (7.37)$$

and 0 otherwise.




Figure 7.2: (a) Histograms of simulated total prior updates U (dots), Eq. (7.6), for the BPM (blue), BT1 (green), and BT2 (red) frameworks. 2^{25} independent bins were simulated for an axion that would produce a mean excess $\mu_a = 5.1$, assuming there is in reality no axion present and using a standard $M = 2$ -identical scan protocol with scan parameters similar to those of HAYSTAC: threshold $x_T = 3.455$ ($f_{n0} = 5\%$, $f_{p0} = 0.03\%$). The error bars were estimated from binomial statistics and are smaller than the data points. The simulated probability distributions are in excellent agreement with the analytic predictions of Eqs. (7.44), (7.35), and (7.32) (blue, green, and red lines). The total updates described in (a) are combined via Eq. (7.7) to form the aggregate prior updates \mathcal{U} described in (b). The aggregate prior updates are numerically generated over 2^{21} independent trials of $N = 2^{15}$ independent bins for each of the three analysis frameworks. For the frameworks for which we derive analytic predictions (green and red lines for BT1, Eq. (7.36), and BT2, Eq. (7.33), respectively), the agreement with the simulations within binomial error bars confirms our expectations. For all three frameworks, the (relatively few) bins which fall outside of the plotted windows do indeed balance the scales so as to produce the unbiased expectation values $\langle U \rangle = \langle \mathcal{U} \rangle = 1$ (Sec. 7.1.4). The BPM framework displays by far the lowest median outcomes in all cases, indicating that it will typically achieve superior exclusion when no axion is present. Figure reproduced from Ref. [126].

For the BPM framework, we restrict ourselves to $M = 2$ identical scans, where the second occurs conditionally on the first exceeding threshold, and derive only the expression for the total

prior update. To derive the PDF for the two-scan BPM prior update, we consider two mutually exclusive cases.

First, with probability $1 - f_{p0}$, the initial scan does not exceed threshold ($x < x_T$; a true negative). In that case, the prior update PDF is

$$f_{u,\text{tn}}(u_{\text{tn}}) = f_x(x) \left(\frac{du}{dx} \right)^{-1} = \begin{cases} (1 - f_{p0})^{-1} f_u(u) & 0 \leq u_{\text{tn}} < u_T \\ 0 & \text{otherwise,} \end{cases} \quad (7.38)$$

where

$$f_u(u) = \begin{cases} \left(\frac{\exp[-\mu_a^2/8]}{\mu_a \sqrt{2\pi}} \right) \left(\frac{\exp[-(\log u)^2/2\mu_a^2]}{u^{3/2}} \right) & 0 \leq u \\ 0 & u < 0 \end{cases} \quad (7.39)$$

and

$$u_T = \exp \left[-\frac{\mu_a^2}{2} + \mu_a x_T \right] \quad (7.40)$$

is the single-scan prior update achieved at threshold ($x = x_T$). The second line of Eq. (7.38) follows from Eqs. (6.15) and (7.4). The normalization factor $(1 - f_{p0})$ accounts for the nonunit probability of drawing from this distribution in the initial scan.

Next we consider the case where the first scan exceeds threshold, ($x \geq x_T$; a false positive), which occurs with probability f_{p0} . The total prior update for the initial false-positive is the product of two single-scan prior updates $U_{\text{fp}} = u_{\text{fp}} u$. The PDF for the first scan prior update u_{fp} is

$$f_{u,\text{fp}}(u_{\text{fp}}) = \begin{cases} (f_{p0})^{-1} f_u(u_{\text{fp}}) & u_T \leq u_{\text{fp}} \\ 0 & u_{\text{fp}} < u_T, \end{cases} \quad (7.41)$$

and that for the second scan prior update u , which is not restricted to being above or below threshold, is given by Eq. (7.39).

The PDF for U_{fp} is then

$$f_{U,\text{fp}}(U_{\text{fp}}) = \int_{-\infty}^{\infty} f_u(u) f_{u,\text{fp}}\left(\frac{U_{\text{fp}}}{u}\right) \frac{du}{|u|} = \begin{cases} \frac{\exp[-(\mu_a/2)^2 - (\log[U_{\text{fp}}/2\mu_a])^2]}{f_{p0}\mu_a U_{\text{fp}}^{3/2} 4\sqrt{\pi}} \operatorname{erfc}\left(\frac{\log[u_T^2/U_{\text{fp}}]}{2\mu_a}\right) & 0 \leq U \\ 0 & U < 0. \end{cases} \quad (7.42)$$

The total prior update U for the procedure is a probability-weighted sum of PDFs:

$$f_U(U) = (1 - f_{p0})f_{u,\text{tn}}(U) + f_{p0}f_{U,\text{fp}}(U). \quad (7.43)$$

Written out fully,

$$f_U(U) = \begin{cases} \frac{\exp[-(\mu_a/2)^2 - (\log[U/2\mu_a])^2]}{\mu_a U^{3/2} 4\sqrt{\pi}} \operatorname{erfc}\left(\frac{\log[u_T^2/U]}{2\mu_a}\right) + \frac{\exp[-\mu_a^2/8 - (\log U)^2/2\mu_a^2]}{\mu_a U^{3/2} \sqrt{2\pi}} & 0 \leq U < u_T \\ \frac{\exp[-(\mu_a/2)^2 - (\log[U/2\mu_a])^2]}{\mu_a U^{3/2} 4\sqrt{\pi}} \operatorname{erfc}\left(\frac{\log[u_T^2/U]}{2\mu_a}\right) & u_T \leq U \\ 0 & U < 0 \end{cases} \quad (7.44)$$

is the two-identical scan BPM total prior update PDF: that is, the probability density for obtaining a prior update U in a single axionless bin using the standard initial and conditional-rescan procedure. The total updates for larger numbers of scans and the aggregate updates for multiple bins are too cumbersome to write out in closed form, but may readily be simulated.

Equations (7.44), (7.35), and (7.32) are respectively the total prior update probability distributions for the BPM framework and the two BT frameworks discussed in Sec. 7.3. These three equations are plotted as blue, green, and red lines, respectively, alongside simulated total prior updates for the three frameworks (blue, green, and red dots) for a two-scan protocol without an axion present in Fig. 7.2a.

The simulations are performed by drawing Gaussian random variables according to Eq. (6.15), with $\mu = 0$, $\sigma = 1$, and performing a second draw as well if $x \geq x_T$. Prior updates are then applied according to Eqs. (7.4), (7.6), (7.27), (7.29), and (7.31). The simulations match the analytic expressions, and together reveal that the vast majority of bins will receive a downward update under all three frameworks. In particular, the total prior updates for BT1 and BT2 almost always take

the precise values f_n and $1 - (1 - f_n)^2$, respectively. The price we pay for these reliably low typical updates are rare false positives with $U \gg 1$ (far off the right side of the plot). The common $U < 1$ and rare $U \gg 1$ outcomes together ensure $\langle U \rangle = 1$, the condition for an unbiased analysis. In contrast, moderate positive and negative prior updates occur much more frequently in the BPM framework, and $\langle U \rangle = 1$ is enforced by a long tail of positive prior updates. The fact that the median of the BPM distribution is lower than that of either threshold distribution is noteworthy, as we will see presently in our discussion of the aggregate prior update.

Figure 7.2b shows the aggregate prior updates \mathcal{U} for the identical two-scan protocol when there are 2^{15} independent bins, not far from the actual number in HAYSTAC [78]. In all cases, the total updates from all bins are averaged according to Eq. (7.7) to obtain the aggregate prior update. All three methods deliver an unbiased aggregate prior update with mean $\langle \mathcal{U} \rangle = 1$, yet the BPM framework typically delivers a stronger exclusion in the absence of an axion, i.e. its *median* aggregate prior update is smaller than in either threshold framework by a factor of two. The improvement is a consequence of the informational advantage of the BPM framework.

The factor-of-two typical improvement in aggregate exclusion illustrated for typical HAYSTAC parameters in Fig. 7.2b yields a scan rate enhancement consistent with the 36% speedup inferred from real data in Fig. 7.3b, absent an axion.¹⁵⁴ If the time for the BT2 framework to achieve the targeted 90% exclusion across the scan window (factor-of-ten reduction in aggregate probability) is T_{90} , the additional factor of two should take approximately $T_{95} = T_{90} \log_{10}(2)$, roughly 30% longer.

7.4 Test using HAYSTAC Phase I dataset

In order to test the BPM framework, it was used to reanalyze the full dataset from Phase I of the HAYSTAC experiment, which includes at several frequencies as many as $m = 3$ scans. Features specific to this dataset are discussed in Ref. [107]. Notably, several frequency bins in the original analysis were discarded for having excess RF power that could with certainty be identified

¹⁵⁴ With an axion present, additional simulations (not shown) reveal its presence in a manner consistent with the interpretation of the global false-positive rate used in frequentist hypothesis testing.

as at least partially nonaxionic in nature. In the BPM reanalysis, we set $U_i = 1$ for these bins. Otherwise, straightforward application of Eqs. (7.6) and (7.7) yield the results in Figs. 7.3a and b, respectively.




Figure 7.3: (a) Reanalysis of HAYSTAC Phase I data as a test of the BPM framework. The standard 90% exclusion line (red) achieved with the FT framework [107] is equivalent to the 10% prior update contour of the BT2 framework discussed in Sec. 7.3.2. The 10% prior update contour (blue) achieved with the BPM analysis constrains more aggressive couplings at nearly all frequencies. The logarithmic color scale indicates the subaggregated prior updates U_s throughout the full two-dimensional parameter space. Darker (lighter) shading corresponds to increased (decreased) probability of an axion of given frequency ν and axion-photon coupling g_γ existing. Frequencies where rescans were performed are marked with a black dot on the threshold exclusion line. The five least unlikely (LU) axion candidates — those whose priors increased most — are marked at their maximum likelihood couplings with orange circles. In the lower panel, corresponding orange bars indicate on the orange logarithmic axis the total posterior probability $P_{\text{LU}}^{(1)}$ within the BPM framework that an axion exists at each of the LU points normalized to the aggregate prior probability \mathcal{P}_a that an axion resides anywhere in the window at that coupling. The LU candidate at $\nu_{\text{LU}}^* = 5.66417255$ GHz, marked in pink, is more than two times higher than the remainder of the top five combined. (b) Aggregate prior update \mathcal{U} (equivalently, exclusion \mathcal{E}) taken as a function of coupling g_γ over the entire HAYSTAC Phase I frequency window for the FT and BT2 frameworks (solid red line), BPM framework (solid blue line), and BPM framework with several adjustments made for a clean comparison with the performance of the thresholding (dashed blue line; see main text). The dashed, gray line, zoomed in upon in the inset, indicates relative $\mathcal{U} = 10\%$ ($\mathcal{E} = 90\%$) scan time normalized to that achieved with the FT and BT2 frameworks [107]. Under typical experimental conditions, the exclusion achievable with the BPM framework would take an estimated 36% longer to attain using thresholding. Figure reproduced from Ref. [126].

7.4.1 Main results

The heat-map of Fig. 7.3a shows the updates to the probability of the axion existing throughout the two-dimensional parameter space. At high couplings ($g_\gamma \geq g_{\gamma,\text{high}} \approx 3g_\gamma^{\text{KSVZ}}$), the probability of an axion is greatly diminished from what it was prior to collecting data. At low couplings ($g_\gamma \leq g_{\gamma,\text{low}} \approx g_\gamma^{\text{KSVZ}}$), virtually no information can be gleaned from the data about the presence or absence of an axion. At intermediate couplings, $g_{\gamma,\text{low}} < g_\gamma < g_{\gamma,\text{high}}$, updates to the probability of an axion being present depend sharply on the specific measured powers $x_i^{(j)}$, and also on the somewhat smoother haloscope sensitivity profile $\eta_i^{(j)}$.

The FT framework has its 90% exclusion line plotted in red in Fig. 7.3a. This corresponds directly¹⁵⁵ to the exclusion plotted in Ref. [107], and equivalently to the 10% prior update contour of the BT2 framework of Sec. 7.3.2. The solid, blue line is the BPM framework's equivalent 10% prior update contour. At almost all frequencies, it constrains more pessimistic couplings than do the threshold frameworks.

The frequency-dependent update (or exclusion) data of Fig. 7.3a has been aggregated in accordance with Eq. (7.7) and Eq. (6.22) for the BPM and FT frameworks, respectively, in Fig. 7.3b. The solid, red line indicates aggregate updated probability (left axis) or exclusion of (right axis) an axion anywhere within the HAYSTAC Phase I frequency window as a function of coupling. The solid, blue line is the updated probability for the BPM framework. Its upward bulge around $2.5g_\gamma^{\text{KSVZ}}$ is largely due to the least unlikely (LU) candidate discussed in Sec. 7.4.2. Had the knowledge from the BPM framework been in hand before the HAYSTAC Phase I decommissioning [42], little additional scan time would have been required to either identify or promptly rule out an axion signal. Therefore, an exclusion curve more representative of the BPM framework for comparison purposes is given by the dashed, blue curve, where the prior update of the foremost LU candidate is reset to unity. Other more minor differences between the dashed and solid blue curves due to particular features of the HAYSTAC Phase I dataset not representative of typical haloscope data

¹⁵⁵ The correspondence comes with the caveat that the data here has been subaggregated into bins of width ≈ 2 MHz, narrower than would be of practical concern to any experimentalist considering a haloscope run at nearby frequencies.

are discussed in Sec. 7.4.3.

The figure of merit of an axion haloscope is the rate R at which it can scan a given frequency window to a specified sensitivity (Sec. 2.5.2). The appropriate comparison to make between two analysis frameworks is therefore the relative amount of time τ_{rel} taken to perform a given scan. The dashed, gray line in Fig. 7.3b represents the time that would have been required to scan any coupling g_γ to the standard 90% exclusion (10% prior update) level using the FT (BT2) framework. The line derives from the

$$\tau_{\text{rel}} \propto g_\gamma^{-4} \quad (7.45)$$

relation from Eq. (2.16) of scan times to couplings, and is normalized to the run time of HAYSTAC Phase I. The inset of Fig. 7.3b shows all three curves in the vicinity of this line. Had the BPM framework been in place to recommend further scanning at ν_{LU}^* , its $\mathcal{U} = 10\%$ aggregate prior update would have taken $\approx 36\%$ longer to achieve with a threshold analysis. In the context of detection efforts that take years to cover unsatisfactorily little parameter space, this is a significant enhancement available at no hardware or operational expense across a broad array of platforms.

7.4.2 Least-unlikely candidates

At the few frequencies where the BPM framework indicates weaker exclusion than the threshold frameworks, it does so with good reason. Most of these LU axion candidate frequencies had excess power above threshold on an initial scan and not far below threshold on a rescan. This can be seen from the spectral coincidence of several of these LU candidates, marked with orange dots at their ML couplings, with the black dots indicating the performance of at least one rescan. The LU candidates that were not rescanned corresponded to initial scan powers coming in just below threshold.

In the bottom panel of Fig. 7.3a, the prior updates for the five LU candidates are plotted normalized to the aggregate prior probability that an axion resides in the HAYSTAC Phase I window. As discussed in Ref. [95], it is not obvious how to distribute one's prior probability of an axion's existence with respect to coupling g_γ . However all five LU candidates have ML

coupling strengths within a factor-of-1.5 range, wherein priors should not differ greatly. Therefore, the *relative* significance of the highest of these LU candidate prior updates (frequency $\nu_{LU}^* = 5.66417255$ GHz) is noteworthy: it is more than two times higher than the remaining LU candidates combined.

It should be emphasized that despite its exceptionally high prior update relative to the rest of the dataset, it is overwhelmingly unlikely that this LU candidate is an axion. A rough attempt at estimating priors (Sec. 7.1.3) puts an optimistic probability of an axion in this bin still below part-in-1,000, with more realistic estimates well below that. Secondly, as discussed in Sec. 7.1.5, the BPM framework is inherently conservative with respect to exclusion in a real dataset, and consequently liberal in its identification of candidates. In particular, nonaxionic RF power excesses caused other spikes in the HAYSTAC Phase I power spectra. Many of these were manually removed from the dataset on the grounds of supplementary evidence (e.g. failure to persist when the auxiliary, weakly-coupled antenna was pulled out of the cavity) [107]. Since this LU candidate came in under threshold on a third and final scan, it was rejected subject to the FT framework. However, the more informative BPM framework suggests that significantly more than any of the other HAYSTAC Phase I frequencies, it merits further interrogation. Generously precluding the possibility of nonaxionic excess power contamination, its probability of containing an axion at its ML coupling is roughly the probability of 40 MHz of nearby, unscanned parameter space, or $P_{LU}^{(1)}/\mathcal{P}_a = 20\%$ of the 200 MHz scan window, containing an axion at that coupling.

7.4.3 Additional features of the Phase I dataset

In addition to the LU candidate at ν_{LU}^* discussed in Sec. 7.4, the dashed, blue curve of Fig. 7.3b also neglects two other features of the HAYSTAC Phase I dataset that appear in Fig. 7.3a. First, the BPM reanalysis includes data from rescans over the upper 100 MHz of the scan window that were performed during Phase I, Run I of HAYSTAC [90, 106]. These rescans were discarded, and new ones performed over some of the same frequencies, because an error in estimating the expected axion lineshape lowered their sensitivities below the levels required by the predetermined confidence

levels and initial scan sensitivities [78]. The extra rescan data itself is valid, however, and is thus included in Fig. 7.3a and b, showcasing the flexibility of BPM to integrate all available information. This extra data is not included in the comparison (dashed blue) curve, as the existence of extraneous rescan data is not ordinary in haloscope analyses.

Secondly, a 3.94σ grand spectrum excess at grand spectrum bin 550 ($\nu_{550} = 5.59886155$) GHz was recorded in the initial scan. This excess occurs 0.3% of the way into the scan window, where the number of contributing raw spectra (Sec. 6.1) was 14, considerably less than the typical 40. The lower number of spectra implies at once an increased susceptibility to systematic error and a reduced axion sensitivity $\eta_{550}^{(1)} = 0.23$. The reduced sensitivity implies that an axion here has its prior update, Eq. (7.4), maximized at $g_\gamma \approx 4.1$. This excess ought to have been rescanned during Phase I of HAYSTAC, but was not. As such, its positive prior update is included in the BPM reanalysis. Its presence is indicated by the dark column at the left of Fig. 7.3a, and is solely responsible for the failure of the aggregate prior update (solid blue line, Fig. 7.3) to vanish at high couplings $\gtrsim 3.2g_\gamma^{\text{KSVZ}}$. Had the required rescans been performed, a source of excess would likely have been discovered or rejected in short order, as in the case of the LU candidate at ν_{LU}^* . As such, it is removed from the comparison (dashed blue) aggregate update curve.

7.5 Concluding remarks on statistical analysis frameworks for axion searches

Haloscopes belong to a class of experiments that are limited by the amount of time practically available to acquire data. As such, optimizing the extraction of the information content of their data is a high priority. Analyzing haloscope data with threshold frameworks (Secs. 6.3, 7.3.1, and 7.3.2), however, discards information pertinent to the presence or absence of an axion, and the alternative frequentist approach of using p -values (Sec. 6.4) is even more problematic. The BPM analysis framework straightforwardly applies Bayes' theorem to incorporate the relevant information content of haloscope power spectra into a posterior probability of an axion existing. Taken in ratio with the prior probability of that axion, the updated probability is seen to be completely equivalent to the standard exclusion quoted in the literature for FT frameworks. Upon applying the BPM framework

to the Phase I data of the HAYSTAC experiment, we observe exclusion improved commensurate with what under typical experimental conditions would translate to a 36% improvement in scan time — within the sensitivity error bounds for the HAYSTAC Phase I dataset, though independent of the contributing sources of error [106]. Additionally, the BPM framework better spotlights those regions of parameter space where the probability of an axion increases. For HAYSTAC’s data, a single such frequency worthy of (but not demanding) further attention stands out.

The BPM framework may be straightforwardly applied to data from ADMX [91], HAYSTAC [107], and other haloscopes and similar dark matter searches [75, 94, 95, 108–111, 113, 116, 244, 251]. BPM will typically probe deeper couplings at no experimental cost while permitting experimentalists to operate nimbly, without the stricter constraints of preordained scan protocols. This latter advantage is especially important for narrowband, tunable searches, in which the question of how to optimize a scan protocol by responding in real-time to acquired data is of paramount importance. Enhancements gained through operational efficiency will readily compound with those intrinsically available from BPM’s informational advantage over threshold-based analyses.

In the following chapter, the BPM framework will be used not in the context of a reanalysis, but on new data from the HAYSTAC experiment. Not only will it increase scan efficiency over FT as predicted in this chapter, but its enhanced operational freedoms (Sec. 7.3), not available in the context of a reanalysis of already existing data, will compound that increase. In Ch. 9, we will see the versatility of BPM as it is applied to a broadband search for ALPs extending as much as 17 orders of magnitude lower in mass than the QCD axions of this chapter and the next.

Chapter 8

A quantum-enhanced search for dark matter axions

In Sec. 2.5.3, we saw the strong need for improving haloscope scan rates in order to sweep out meaningful portions of parameter space on acceptable timescales. I have argued that the most promising advances are those that either (1) can be iteratively improved upon in order to continue the decades-long trend of steadily increasing scan rates or (2) are easily and broadly applicable at little or no cost to the rest of the experiment.

In Ch. 5, we saw a direct demonstration of the squeezed state receiver (SSR) concept [208], which achieved a factor-of-two scan rate enhancement in a realistic mock-haloscope experiment. In line with objective (1) above, the SSR and quantum technologies for axion searches more generally also promise scalability as transmission losses continue to come down [149, 218, 222] or reciprocal elements are eliminated [252]. The history of improvements to haloscopic scan rate, wherein the vast majority of improvement over the past three decades has come from reducing noise as opposed to increasing signal, further recommend the SSR approach.

In Ch. 7, we discussed an entirely separate and complimentary advance in the spirit of objective (2). There, the Bayesian power measured (BPM) analysis framework [126] improved the scan efficiency of the HAYSTAC Phase I analysis (Sec. 7.4) by what would have been, under normal operating conditions, $\approx 36\%$. While improved data handling and related scan strategy considerations are of secondary importance to fundamentally altering the noise regime in which modern haloscopes operate, the advantage of improved data analysis is that it comes at zero hardware and *negative* operational cost (Sec. 7.3).

This chapter, which presents the work of [253], combines the advances of Chs. 5 and 7 to present the most important result of this thesis: a live haloscope run directly improved by the SSR apparatus and the BPM framework. We begin in Sec. 8.1 with the experimental setup and operation of the haloscope. Section 8.2 takes a detailed look into the data processing for the experiment, and Sec. 8.3 presents analysis and results. Finally, Sec. 8.4 covers calibration measurements taken periodically throughout the run to determine the noise level of the experiment. The combined benefit of SSR and BPM amounts to a tripling of the scan rate for dark matter axions in Phase II¹⁵⁶ of the HAYSTAC experiment over a mass range favored by recent theoretical projections [69, 70]. The data run reveals no signature of dark matter axions in the combined 16.96–17.12 and 17.14–17.28 $\mu\text{eV}/c^2$ mass window for axion-photon couplings above $g_\gamma = 1.38g_\gamma^{\text{KSVZ}}$, reporting exclusion at the 90% level. The use of quantum measurement in particular in an axion search marks the first time quantum squeezing has been harnessed to improve the sensitivity of a search for fundamental particles, inviting a new era of quantum-enhanced searches for physical phenomena.

8.1 Setup and operation

Phase II of the HAYSTAC experiment is shown schematically in Fig. 8.1. The axion cavity shown in Fig. 8.2 resides in an LD250 BlueFors dilution refrigerator operated at a base temperature of 61 mK and shown in Fig. 8.3. A solenoidal, superconducting magnet sources a spatially homogeneous $B_0 = 8$ T field. The cavity sits at the center of the high magnetic field region, Fig. 8.4, and has volume $V_C = 1.545$ L (not counting space occupied by the tuning rod)¹⁵⁷ and typical $Q_0 = 47000 \pm 5000$ [96].

Two Josephson parametric amplifiers (JPAs; SQ and AMP, as in Chs. 4 and 5) hang from a lower plate thermalized to the mixing chamber plate (Fig. 8.5a) approximately 1 m above the cavity in a “field-free” region engineered by a counterwound bucking coil integrated into the magnet. The

¹⁵⁶ The HAYSTAC data run discussed in this thesis is more properly Run I of Phase II. For brevity, I will leave off the run designation herein, though at least one future Phase II run is planned.

¹⁵⁷ Alternatively, the cavity could be treated as a 2 L volume with a commensurate decrease to the form factor, Eq. (2.15).

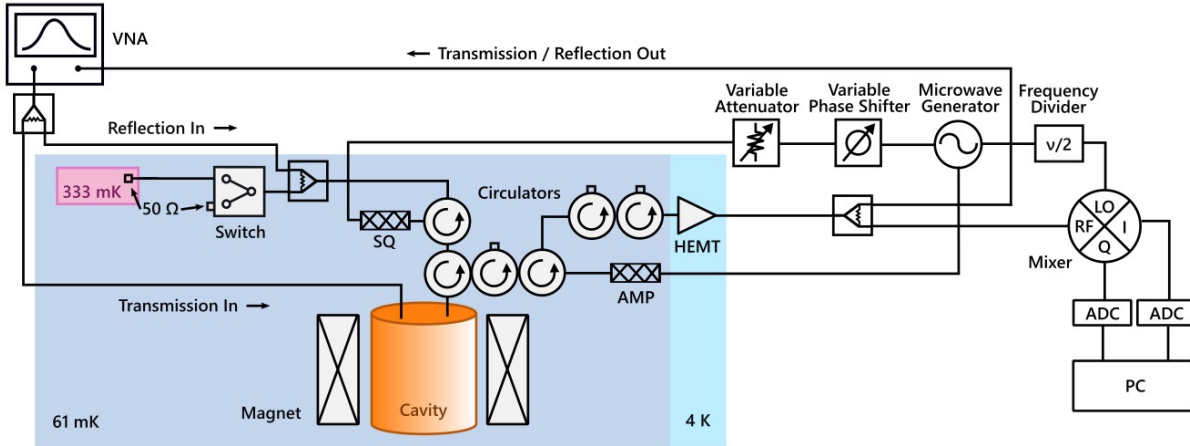


Figure 8.1: Simplified HAYSTAC Phase II experimental schematic. A single signal generator provides the local oscillator (LO) tone as well as the tones for pumping both JPAs. Each Josephson parametric amplifier (JPA) has two ports: one for the input of probe tones and one where the signal is input/output. The LO is set at half the pump frequencies via a frequency divider, and the relative phase and amplitude of the pump tones are set using a variable phase shifter and attenuator on the squeezer (SQ) pump line. Switches in the SQ and amplifier (AMP) pump lines (not shown) are used to toggle the JPAs on and off. Microwave circulators route signals nonreciprocally in order to realize the time sequence of operations used to benefit from squeezing (Sec. 5.4). Circulators with a $50\ \Omega$ termination on one port act as isolators, shielding upstream circuit elements from unwanted noise coming from further down the measurement chain. During data acquisition and calibration measurements, signal and noise emitted from and reflected off the cavity are amplified by a high-electron-mobility transistor (HEMT) amplifier at 4 K, fed into the radio frequency (RF) port of an in-phase/quadrature (IQ) mixer and mixed down to an intermediate frequency (IF), digitized (analog-to-digital converter; ADC), and recorded (personal computer; PC) where the power spectral density is calculated. The cavity's Lorentzian profile is monitored with reflection and transmission measurements taken using a vector network analyzer (VNA), for which a portion of the output is split off before the IQ mixer. A switch that toggles between a hot (333 mK) and cold (61 mK) $50\ \Omega$ loads is used for calibration measurements described in the text. Figure reproduced from Ref. [253].

JPAs are further shielded by a 4-layer shielding can (niobium, Amumetal 4K, aluminum, Amumetal 4K), surrounded by three superconducting bucking coils (Fig. 8.5b). As demonstrated in Fig. 8.6, this shielding is sufficient to reduce the solenoid-sourced flux through the JPAs' SQUID loops to much less than one magnetic flux quantum, a roughly millionfold decrease in field strength from the nearby high-field region. The cavity is tuned by rotating the copper-plated tuning rod off-axis (Figs. 2.6b and 8.2).

At each operating point, The TM_{010} -like mode is stepped using an ANR240 piezoelectric




Figure 8.2: HAYSTAC cavity photo. The copper-coated stainless steel cylindrical cavity [96] has height 10 in and radius 4 in. The axion-sensitive TM_{010} -like mode is tuned by rotating the 1 in tuning rod off-axis. The empty cavity has volume 2 L, which reduces to $V = 1.545$ L with the rod present.

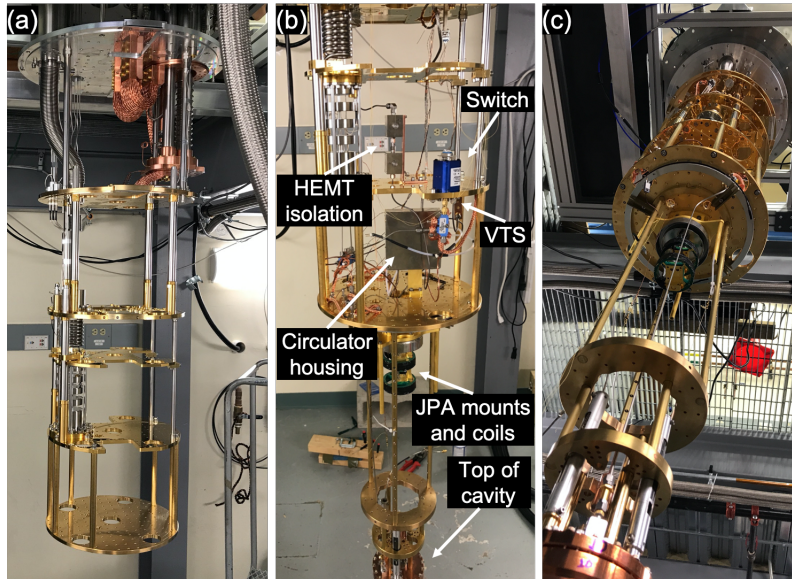


Figure 8.3: (a) Photo of the LD250 BlueFors dilution refrigerator housed at Yale before cabling and haloscope components were installed. (b), (c) Side- and angled-view photos where key elements labeled in (b) have been installed.

motor to rotate the tuning rod, whereupon a vector network analyzer is used to extract the cavity resonance ν_c , loaded cavity quality factor¹⁵⁸

$$Q_L = \frac{2\pi\nu_c}{\kappa_m + \kappa_l}, \quad (8.1)$$

¹⁵⁸ The coupling rates κ are as defined in Ch. 5.

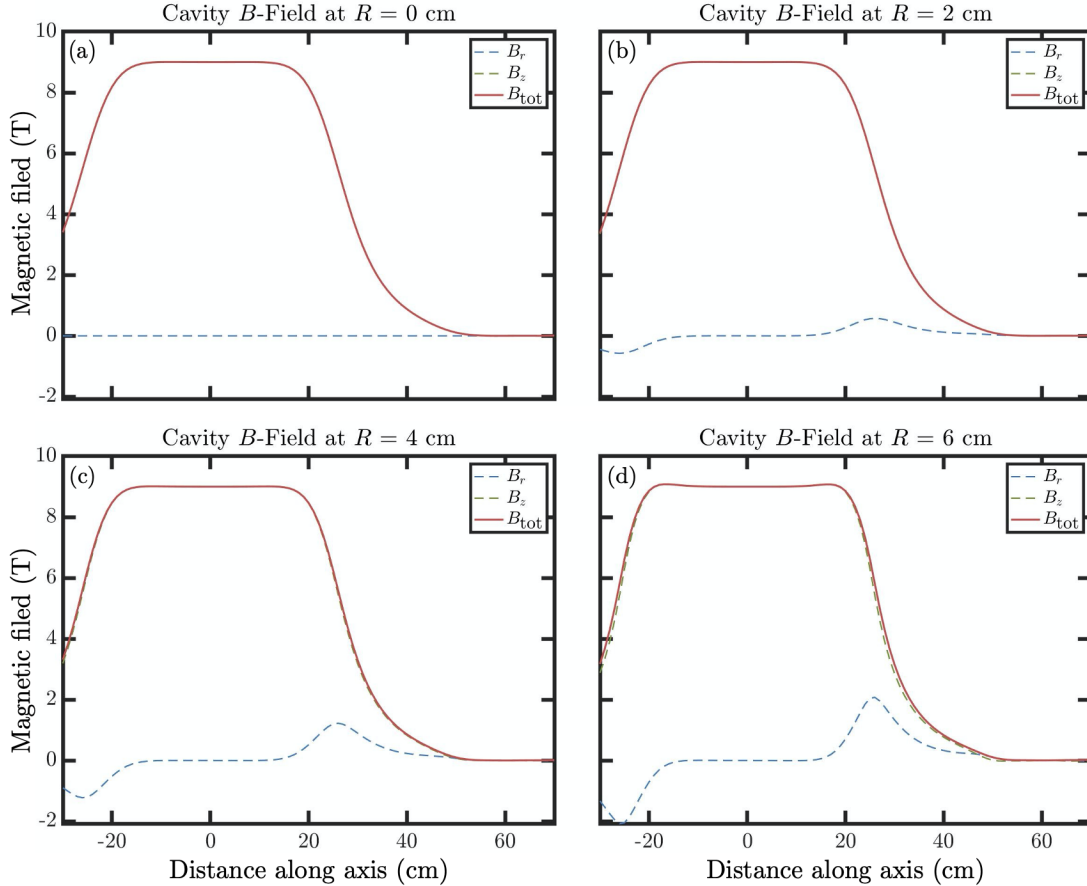


Figure 8.4: (a)–(d) Radial (B_r), axial (B_z), and total (B_{tot}) magnetic field specified values at maximum field strength and $R = 0$ (a), 2 (b), 4 (c), and 6 (d) cm distances from cavity axis as a function of distance along the cavity axis for the solenoid used in HAYSTAC Phases I and II. A change in field component sign indicates a reversal of direction. The dominant contribution over the entire region of the cavity of height 10 in and radius 4 in comes from B_z , the component useful for axion detection. The maximum field had to be lowered from 9 to 8 T for Phase II due to damage sustained in magnetic quenches during Phase I operation [42, 90].

and coupling factor

$$\beta = \frac{\kappa_m}{\kappa_l} \quad (8.2)$$

via pairs of transmission and reflection measurements. In order to center the amplification and analysis bands on cavity resonance, a single microwave generator sets the two JPA pump frequencies to $\nu_p = 2\nu_c$ and the same tone is further split off and frequency-divided to become the local oscillator (LO) used for homodyne measurement about ν_c . The relative pump phase θ and ampli-




Figure 8.5: (a) Mixing chamber stage layout shown in SolidWorks. The gold-plated mixing chamber plate is thermalized to a gold-plated lower plate from which hangs the JPA housing and below that (not shown) the cavity. The JPAs are placed in the nominally “field-free” region, and the three gray rings show the locations of static superconducting bucking coils which are used to cancel the remaining field through Lenz’s law. The four circulators required to implement the SSR in practice (Fig. 5.8) are housed in magnetic shielding above the lower plate. (b) Photo of the JPA magnetic shielding. The static bucking coils sit outside a four-layer magnetic shield; from inside-out: niobium, Amumetal 4K, aluminum, and Amumetal 4K (visible layer). The Amumetal 4K, from Amuneal Manufacturing Corp., is a high permeability material around the temperatures where the aluminum and niobium undergo their superconducting phase transitions, so as to minimize flux vortices in the superconducting layers when at base temperature. The whole configuration reduces the 8 T magnetic field 1 m below to less than one part in 10^6 of its peak value (Fig. 8.6).

tude are optimized to maximize squeezing using an electronically actuated variable phase shifter and attenuator. Axion-sensitive voltage fluctuations $V_X(t)$ are collected for $\tau = 3600$ s before being Fourier transformed and having their power spectral densities (PSDs) computed.

As in Ch. 5, the SQ and AMP are coupled to the axion cavity in an SSR configuration via a microwave circulator. The SQ prepares a squeezed vacuum state, which is coupled into the axion cavity and subsequently measured noiselessly using the AMP. The setup delivers 4.0 dB of off-resonant vacuum squeezing *after* the state is degraded by transmission losses and added noise (Footnote 90), yielding a 1.9-fold¹⁵⁹ scan rate enhancement beyond what could be achieved at the

¹⁵⁹ This improvement comes relative to the optimal unsqueezed case for the system accounting for improper tuning rod thermalization (Appx. C) in both cases. It also accounts for our measured added noise and the fact that the homodyne measurement leaves a small range of frequencies near DC too noisy to use in both cases. In theory, these lost frequencies imply a higher overcoupling for optimal performance both with and without squeezing, though the




Figure 8.6: Measurement of JPA tunability with magnetic field on and off. One of the flux-pumped JPAs is tuned by a DC current source with the large magnetic field used to search for axions at 0 and 8.4 T. Tuning the magnetic field on causes only a very minimal shift in the tuning structure of the JPA, confirming the proper functioning of the magnetic shielding.

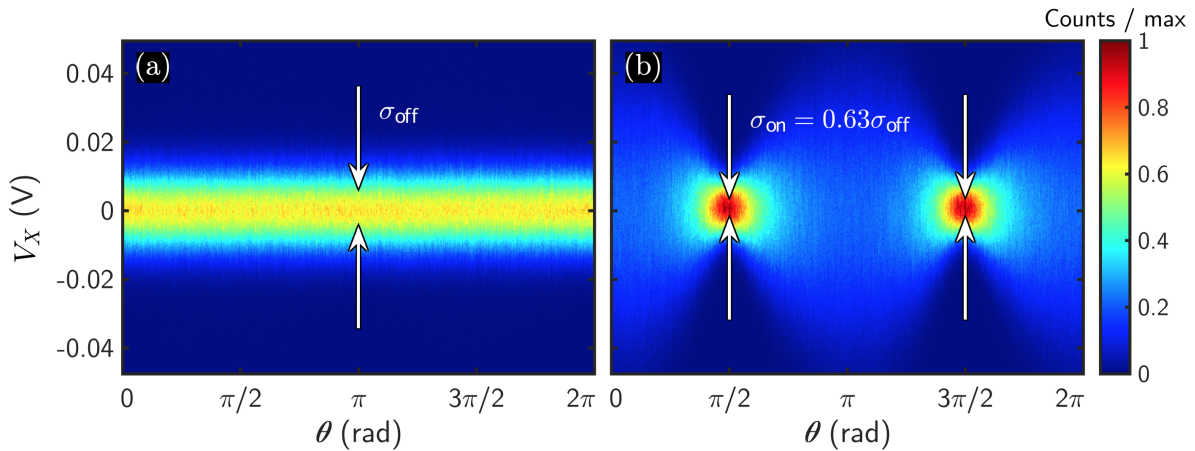


Figure 8.7: (a)–(b) Voltage fluctuation (V_X) histograms as a function of the phase θ between the SQ and AMP pumps first with the SQ off (a) so that it acts like a mirror and there is no preferred phase. With SQ on (b), the variance is minimized for $\theta = \pi/2$ and $3\pi/2$. Figure reproduced from Ref. [253].

effect is slight, and neglected here for both cases. The calculation assumes an analysis band set to a fixed number of (overcoupling-dependent) cavity linewidths.

quantum limit. This measured squeezing,¹⁶⁰ shown in Fig. 8.7, is consistent with the loss and added noise measurements discussed in Sec. 8.4.




Figure 8.8: Illustration of the squeezed state receiver-equipped haloscope showing the transformation of the vacuum state in quadrature space. A vacuum state, whose Wigner function (color surface) [254] is symmetric in quadratures \hat{X} and \hat{Y} , is sourced as Johnson-Nyquist noise from a $50\ \Omega$ microwave termination (black box) at 61 mK. It is routed by a nonreciprocal element (Circ) to the SQ JPA which squeezes the \hat{X} quadrature. The squeezed state may then be displaced by a hypothetical axion field a (red lines) in the axion cavity (Cav). It is subsequently unsqueezed by the AMP JPA, in the process amplifying the axion-induced displacement. The resulting state is read out by a conventional readout chain led by a HEMT amplifier and digitized (spectrum analyzer; SA). Figure reproduced from Ref. [253].

The data acquisition procedure at each tuning step in the initial scan¹⁶¹ is essentially that of the mock-haloscope experiment discussed in Sec. 5.4, except here the cavity really is tuned in search of the hypothesized axion signal, and the live acquisition time is brought up from 0.32 s to $\tau = 3600$ s as in Sec. 8.1. Figure 8.8 illustrates the journey of the squeezed state in quadrature space as it propagates through the SSR-equipped haloscope. The JPAs in this picture transform phase space, turning the vacuum state into a squeezed state and then back into a vacuum state.¹⁶² An

¹⁶⁰ The measurement protocol is the same as that for Fig. 5.5, here showing SQ off in addition to SQ on data.

¹⁶¹ Because squeezing does not improve the maximum single-frequency sensitivity of the haloscope, it is not used during rescans. In theory, a marginal benefit could have been achieved for adjacent bins by squeezing even while critically coupled (Sec. 5.6).

¹⁶² In practice, the gain of the AMP is turned up higher than that of the SQ, so the final state is really elliptical along the opposing axis to the initial squeezed state. For illustrative purposes, however, it is helpful to think of the AMP as simply undoing the transformation of the SQ.

axion, if present, imparts a displacement, and the portion of that displacement along the squeezed quadrature stands out more strongly from the noise than in the case without squeezing.¹⁶³




Figure 8.9: Advantage conferred by squeezing calculated with parameters used in HAYSTAC. Blue curves show Johnson-Nyquist noise S_r reflected off of the cavity as a function of detuning δ_ν from cavity resonance. Red curves show cavity Johnson-Nyquist noise S_c which originates in the cavity. Solid (dashed) lines indicate HAYSTAC Phase II (Phase I) operation. The combination of squeezing and higher cavity coupling increases the bandwidth over which the haloscope is sensitive to axions. Figure reproduced from Ref. [253].

As we saw in Ch. 5, squeezing improves the bandwidth over which the apparatus is sensitive to an axion, rather than the peak sensitivity obtained at cavity resonance. In the setup of Fig. 8.1, the reflected noise is sourced from a 50Ω termination held at base temperature 61 mK, which dominates away from cavity resonance. The ratio of the signal power that would be delivered by an axion at any given frequency to cavity Johnson-Nyquist noise S_c at that frequency is spectrally constant. Because squeezing selectively removes reflected noise S_r , it is beneficial to overcouple the cavity's measurement port by setting β , Eq. (8.2), to twice the deliverable squeezing S , Eq. (4.4), as in Ch. 5. In Fig. 8.9, the dashed lines show cavity thermal Johnson-Nyquist in red and reflected noise in blue with $\beta = 2$, the ideal value for a quantum-limited haloscope. Cavity noise, and therefore a hypothetical axion signal, is only dominant over the narrow band shown by the upper black

¹⁶³ This coherent picture of the axion displacement does a good job showing what each individual realization of the axion field does. Figure 5.1 presents a complimentary picture: there, many coherent displacements of the sort shown in Fig. 8.8 average together to add variance to the squeezed or vacuum state over the course of a measurement. In both pictures, the essential conclusion is that the axion is disproportionately more visible along the squeezed quadrature.

arrow. The solid lines show the operating conditions for this work, where squeezing is employed and the cavity is overcoupled at $\beta = 7.1$.¹⁶⁴ Squeezing reduces the reflected noise by 4.0 dB, and increasing β to 7.1 increases the cavity bandwidth. Without squeezing, $\beta = 7.1$ would make cavity reflected noise dominant at all frequencies. Squeezing this noise below the quantum limit yields near-maximal visibility extended far off resonance.

8.2 Data processing and analysis

The HAYSTAC Phase II data run used the SSR apparatus of Sec. 8.1 to probe over 70 MHz of well-motivated parameter space [69, 70] in half the time that would have been required for unsqueezed operation, saving approximately 100 days of scanning. Initial data acquisition occurred from September 3 to December 17, 2019, covering 4.100–4.178 GHz and skipping a crossing with a transverse electric (TE) mode at 4.140–4.145 GHz which does not couple to the axion [96]. A total of 861 spectra were collected as part of the initial scan, of which 33 were cut due to cavity frequency drift, poor JPA performance, or an anomalous power measurement in a probe tone injected near cavity resonance. The cuts made on the initial spectra are shown in Fig. 8.10. Analysis of the initial scan data yielded 32 power excesses that merited further scanning, consistent with statistical expectations. 508 rescan spectra were collected from February 25 to April 11, 2020, with 3 cut. None of the power excesses from the initial scan persisted in the analysis of rescan data. Calibration measurements taken at each tuning step are described in Sec. 8.4.

The data processing for an axion haloscope, described in Sec. 6.1 and based upon Ref. [78], is a complicated procedure prone to human error and confirmation bias. Furthermore, it requires modifications (Sec. 5.7) before being applied to a system performing single quadrature measurement and using squeezing. To mitigate human error, HAYSTAC used two separate and semi-independent¹⁶⁵ processing and analysis teams, led by Kelly Backes at Yale and myself at JILA, respectively. The

¹⁶⁴ The overcoupling here is slightly higher than twice the delivered squeezing S because the cavity noise is higher than the temperature of the fridge (Sec. 8.4). When not squeezing, the optimal overcoupling is actually $\beta \approx 2.8$, slightly higher than the value $\beta = 2$ plotted in Fig. 8.9. Fully accounted for, the cavity noise has negligible affect on the scan rate enhancement achieved via squeezing. For a quantitative discussion, see Appx. C.

¹⁶⁵ Each team was never allowed to look at the other’s processing and analysis code, but other communication was




Figure 8.10: Raw spectrum cuts represented by black, vertical lines for the 861 raw spectra (horizontal axis, in chronological order) in the HAYSTAC Phase II initial scan. Spectrum cuts were applied for spectra with unsuccessful squeezing (“noSQ”), cavity resonance drift as measured in reflection (“ $\Delta\nu_{\text{rfl}}$ ”) or transmission (“ $\Delta\nu_{\text{tx}}$ ”), anomalous average subspectrum power (“ P_{spec} ”), or anomalous average (“ P_{tone} ”) or standard deviation (“ $\sigma_{P_{\text{tone}}}$ ”) of heights of a probe tone situated alongside the spectra and measured throughout their acquisition. Any one cut condition being met (“any”) was deemed sufficient to cut a spectrum. In total, 33 cuts were made on the initial spectra, and 3 on the rescan spectra (not shown).

two teams ultimately agreed in the final prior updates $U_i(g_\gamma)$ (Sec. 7.1.1) to better than 1%, with the official results reported in Sec. 8.2 being the update-space average of the analyses.

This section follows the JILA processing and analysis from data acquisition to exclusion, presenting a comparison to the frequentist threshold (FT) analysis used in Phase I, Run I [106]. The improved informational efficiency and greater operational freedom of the BPM framework together improve scan rate by 52%. Unlike for the Phase I data of Sec. 7.4, the final results are remarkably clean, with no prominent least unlikely (LU) candidates of note.

permitted.

8.2.1 Subspectra and raw spectra

The 861 initial scan plus 508 rescan raw spectra¹⁶⁶ are each acquired at a sample rate of 10 MS/s in 5 s acquisitions. These 5 s acquisitions are divided in code into 500 separate 10 ms chunks, which are Fourier transformed at $\Delta_b = 1/(10 \text{ ms}) = 100 \text{ Hz}$ Fourier bin resolution and rotated into the AMP's amplified quadrature prior to having their PSDs computed per Step 3 of Sec. 6.1.¹⁶⁷ A 1 hr raw spectrum acquisition therefore collects and averages 360,000 subspectra.¹⁶⁸ To save disk space, only the raw spectra themselves are recorded.¹⁶⁹

Figure 8.11 shows raw spectra acquired in the initial scan. The representative spectrum in 8.11a was randomly chosen as spectrum no. 702. In order to follow the processing of a single spectrum, the same spectrum is in several subsequent plots throughout this section. The irreversible ambiguity between the AMP's signal and idler inputs means that this and subsequent individual spectra are mapped symmetrically about the LO (Sec. 5.7). The gap in the middle of this and all other spectra comes from low-frequency noise present in our homodyne measurement.¹⁷⁰ The predominant feature in the raw spectra is the large-scale structure owing primarily to frequency-dependent gains in the amplification chain and the low-pass filtering of the data before digitization. This broad spectral structure must be filtered out per Step 4 of Sec. 6.1.

In order to filter out large-scale structure, the raw spectra are grouped into three batches. This is necessary because, over the course of the months-long acquisition, room temperature electronics were reconfigured to improve performance. These reconfigurations in turn changed the

¹⁶⁶ Some of these occurred to fill in small gaps in the initial scan, and therefore could alternatively be thought of as belonging instead to that group. Since they were acquired with the rescan data, we tally them there.

¹⁶⁷ For phase-sensitive measurement, Step 2 of Sec. 6.1 is replaced with the rotation of the measured quadratures into the amplified and deamplified quadratures, the latter of which is discarded. Step 1 is discussed in Footnote 168.

¹⁶⁸ The Fourier domain quadrature voltages are complex, so in practice the real and imaginary parts have their PSDs separately computed and summed per Step 1 of Sec. 6.1. There are thus 720,000 independently measured powers being averaged at each 100 Hz Fourier bin.

¹⁶⁹ If the full timestream dataset for the above procedure were saved, it would occupy roughly 800 TB: not an impossible number, but an inconvenient one, especially given that the raw spectra in principle contain the full axion-pertinent information content of the measurement per Sec. 7.2.

¹⁷⁰ The lost low-frequency data is a practical disadvantage of the phase-sensitive measurement relative to the phase-insensitive one. It is counterbalanced by the fact that phase-insensitive measurement operates off resonance, where lower JPA gain gives rise to higher added noise. Unlike for the phase-sensitive versus -insensitive effects considered in Sec. 5.1.3, there is no fundamental reason for these two effects to cancel out. In practice, it is reassuring that neither rises to the level of even the smaller ($\sqrt{2}$) impacts on scan rate discussed in Sec. 5.1.3, and that they work in opposite directions.




Figure 8.11: (a) Individual raw spectrum no. 702 plotted in the RF band along with the Savitzky-Golay (SG) filter for its batch of spectra (batch 3). The SG filter shown is for removing structure common to all spectra within the batch. A subsequent filtering step, shown in Fig. 8.12a for this spectrum, will remove the remaining structure. The spectrum is symmetric about its center as a consequence of the phase-sensitive amplification. (b) All raw spectra shown together in the IF band. Batches are determined by changes made to the measurement chain throughout the months-long initial scan data acquisition. IF/RF subscripts are applied to frequency axes in these and other plots for emphasis.

large-scale spectral structure (or “spectral baselines”) of the raw spectra.¹⁷¹ All 861 raw spectra are shown together in Fig. 8.11, where the differences between batch 1 and batches 2 and 3, which have similar baselines, are evident. Each baseline is divided by a Savitzky-Golay (SG) filter ($d = 10$, $W = 500$) applied to its entire batch in the intermediate frequency (IF) band. After the filter is applied, bins which average in the IF (where an axion signal will not line up between different spectra) to a power excess above 4.5σ are cut along with their close neighbors for all spectra within a batch.

8.2.2 Processed spectra

A second SG filter ($d = 4$, $W = 500$)¹⁷² is applied to each spectrum individually in the radio frequency (RF) band to remove the remaining large-scale structure which remains dominant

¹⁷¹ The three batches are each chronologically contiguous. The first contains 201 spectra, the second 74, and the third 586.

¹⁷² For rescans, the second SG filter has $d = 6$, following Ref. [78].




Figure 8.12: (a)–(d) Individualized SG filtering. The raw spectra P_{RF} , normalized by the collective IF SG filters for their respective batches $\bar{P}_{\text{IF}}^{\text{filt}}$ and with 1 subtracted out, are plotted for four different raw spectra. Also plotted are their individualized SG filters. The structure of the individual spectra all look different from one another because common structure was removed in the first SG filtering stage. The individualized SG filters will be divided out to obtain the processed spectra of Sec. 8.2.2.

over the statistical fluctuations (plus possible axion-induced structure). These SG filters, shown alongside four IF SG-filtered raw spectra in Fig. 8.12, look very different for each individual spectrum even within a batch (compare Fig. 8.12a and Fig. 8.12d, both in batch 3). The dissimilarities among individual spectra indicate that the IF SG filtering of Sec. 8.2.1 successfully removed shared structure.

Upon dividing out the individualized SG filters, we are left with the processed spectra. Figure 8.13a shows processed spectrum no. 702, which we have been following through the data processing.

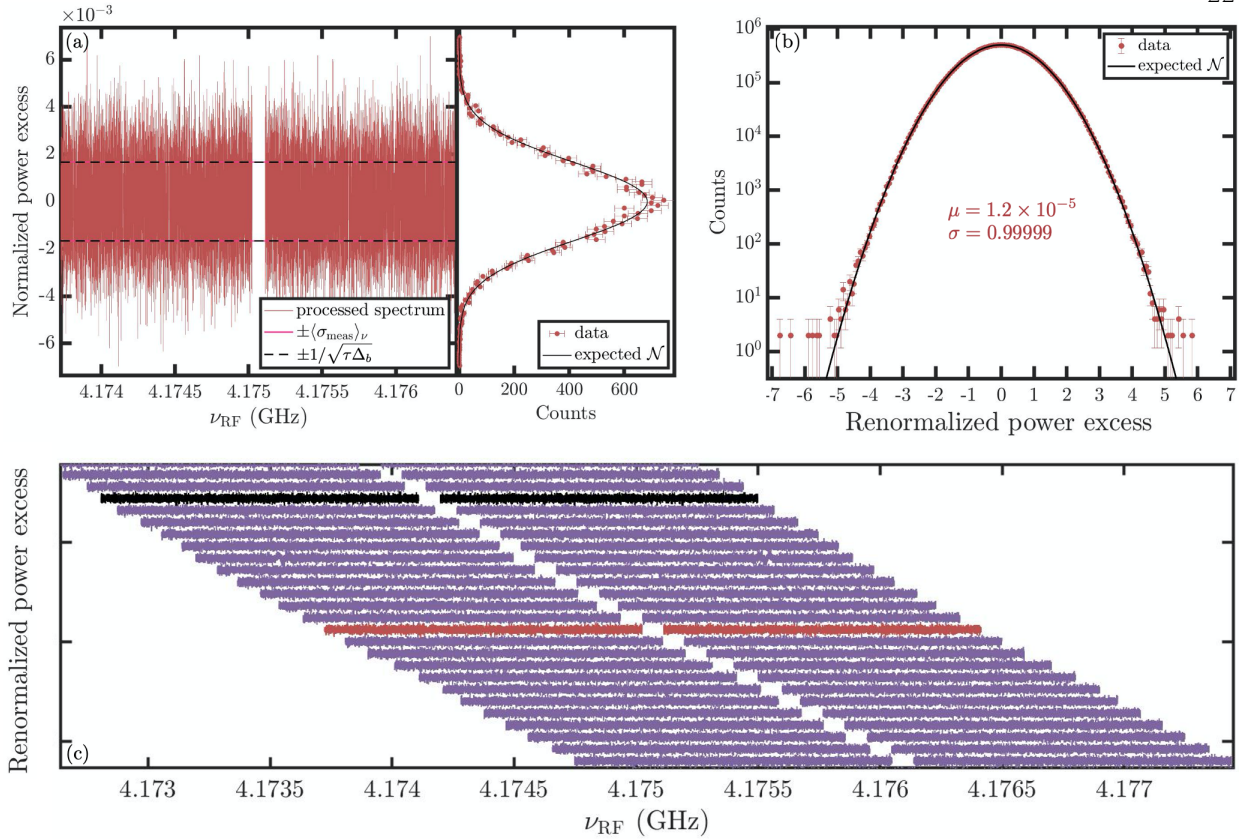


Figure 8.13: (a) Individual processed spectrum no. 702 together with histogram of normalized power excesses. The frequency-averaged inter-processed spectrum standard deviation $\langle \sigma_{\text{meas}} \rangle_{\nu}$ (pink line) agrees well with the expected standard deviation $1/\sqrt{\tau \Delta_b}$ (black, dashed line), and the histogrammed power excess counts for spectrum no. 702 are seen to follow the corresponding expected normal (\mathcal{N}) distribution. (b) Histogram of all renormalized processed spectra obtained by dividing the normalized power excesses of (a) by $\langle \sigma_{\text{meas}} \rangle_{\nu}$. The distribution of all processed spectra is well approximated by the expected Gaussian distribution, and has mean 0, standard deviation 1 as expected. (c) Waterfall plot of many processed spectra near spectrum no. 702 (red) lined up in the RF band. The black spectrum was one of the 33 spectra cut from the initial scan (Fig. 8.10).

It agrees very well with the expected mean 0, standard deviation $1/\sqrt{\tau \Delta_b} = 0.0017$ Gaussian,¹⁷³ as does the collection of all processed spectra excesses shown in Fig. 8.13b. If noise were not averaging down properly in the experiment, the Gaussians would have larger-than-expected standard deviations. In Fig. 8.13c, many processed spectra are shown together, with no. 702 highlighted in

¹⁷³ There are several subtle factors of two that go into Fig. 8.13. First, the standard deviation of $1/\sqrt{\tau \Delta_b}$ could really be written as $(1/\sqrt{\tau \Delta_b}) \times \sqrt{2}/\sqrt{2}$, where the multiplicative $\sqrt{2}$ comes from the standard deviation of the first χ^2 distribution and the divisive factor of two comes from the hidden factor of two in the number of independent contributions to each raw spectrum (Footnote 168). Separately, the error bars in the histograms in Fig. 8.13a and b are $\sqrt{2 \times \text{counts}}$, where the $\sqrt{2}$ multiplies the Poisson standard deviation because of the duplication of all measured powers about the LO by the AMP (Sec. 8.2.1).

red and a cut spectrum in black. Observe that the spectral step size between processed spectra is substantially smaller than the spectra themselves are, which ensures smooth coverage of the scanned frequency window (Sec. 7.1.4). Bins exceeding 6σ are cut from the processed spectra per the procedure in Ref. [78].

8.2.3 Rescaled spectra




Figure 8.14: (a) Rescaling visibility (red line) for spectrum no. 702. Also shown are what the rescaling would be if the AMP gain $G_A(\nu)$ maintained its maximum value of $G_A(\nu_p/2)$ across the entire analysis band (black, dashed line) and if in addition the cavity did not replace any reflected, squeezed noise with internal, unsqueezed noise ($|\chi_{ml}(\nu)|^2 = 0$, black, dotted line). (b) Rescaled spectrum no. 702, obtained from dividing processed spectrum no. 702, Fig. 8.13a, by the rescaling visibility in (a). The rescaled spectra have the special property that bin containing 100% of the power of a KSVZ axion would have mean 1.

The processed spectra are next each rescaled to account for their sensitivity to the axion. Specifically, the single-bin KSVZ axion visibilities $\alpha_{1\text{-bin}}^{\text{KSVZ}}$, shown in red in Fig. 8.14a for spectrum no. 702, are divided out from the processed spectra to produce the rescaled spectra. The rescaling that would have been applied if the AMP maintained its maximum gain across the entire analysis band is shown as a black, dashed line. The fact that the black, dashed line and the red line nearly coincide indicates that JPA bandwidth is not at present a limiting factor, even with the cavity being strongly overcoupled to improve bandwidth upon squeezing. The black, dotted line,

shows the rescaling that would be applied if the cavity still sourced the same axion signal, but without dissipating any of the external noise and replacing it with unsqueezed vacuum. As such, its deviation from the red line, largest towards the center, indicates how much of the squeezed state is being replaced with cavity fluctuations. It shows that we are accounting for the fact that not all of the squeezed state (independent of transmission losses) makes it to the AMP.

Rescaled spectrum no. 702 is shown in Fig. 8.14b. Recall from Step 5 of Sec. 6.1 that the rescaled spectra have the special property that a KSVZ axion whose power falls entirely into a given 100 Hz bin would shift the bin's expectation value from 0 to 1. Thus, bins closer to the center of each raw spectrum, where axion transmission through the cavity is strongest, should be the most sensitive (lowest rescaled standard deviation), as is the case in Fig. 8.14b. The reason for rescaling the spectra to have uniform mean in the presence of an axion, with the sensitivity of each bin now encoded in its standard deviation, is to facilitate the maximum likelihood (ML) estimation used to combine individual spectra.

8.2.4 Combined spectrum

The rescaled spectra are summed with ML weights in order to form the combined spectrum as in Step 6 of Sec. 6.1. The combined spectrum, shown in Fig. 8.15a, has the same property as the rescaled spectrum: a single-bin, KSVZ axion has mean 1. Near the edges of the scan, the expected standard deviation obtained from Eq. (6.7) grows very large, indicating insensitivity to the axion. This is also true around the TE mode crossing near the center of the scan window. Mainly, the sensitivity profile is determined by the number of contributing raw spectra (shown in gray on the right axis) to each combined spectrum bin.

The combined spectrum can alternatively be viewed as in Fig. 8.15b, such that each bin has standard deviation 1 regardless of the axion-sensitivity (and whether or not there is an axion). The sensitivity to the axion is then encoded in the mean of the axion distribution for each bin (not shown). The representation of Fig. 8.15b is better for examining the statistics of the combined spectrum. We find that its measured standard deviation matches the expected value, and it is well

described by the expected Gaussian distribution plotted in the right panel.

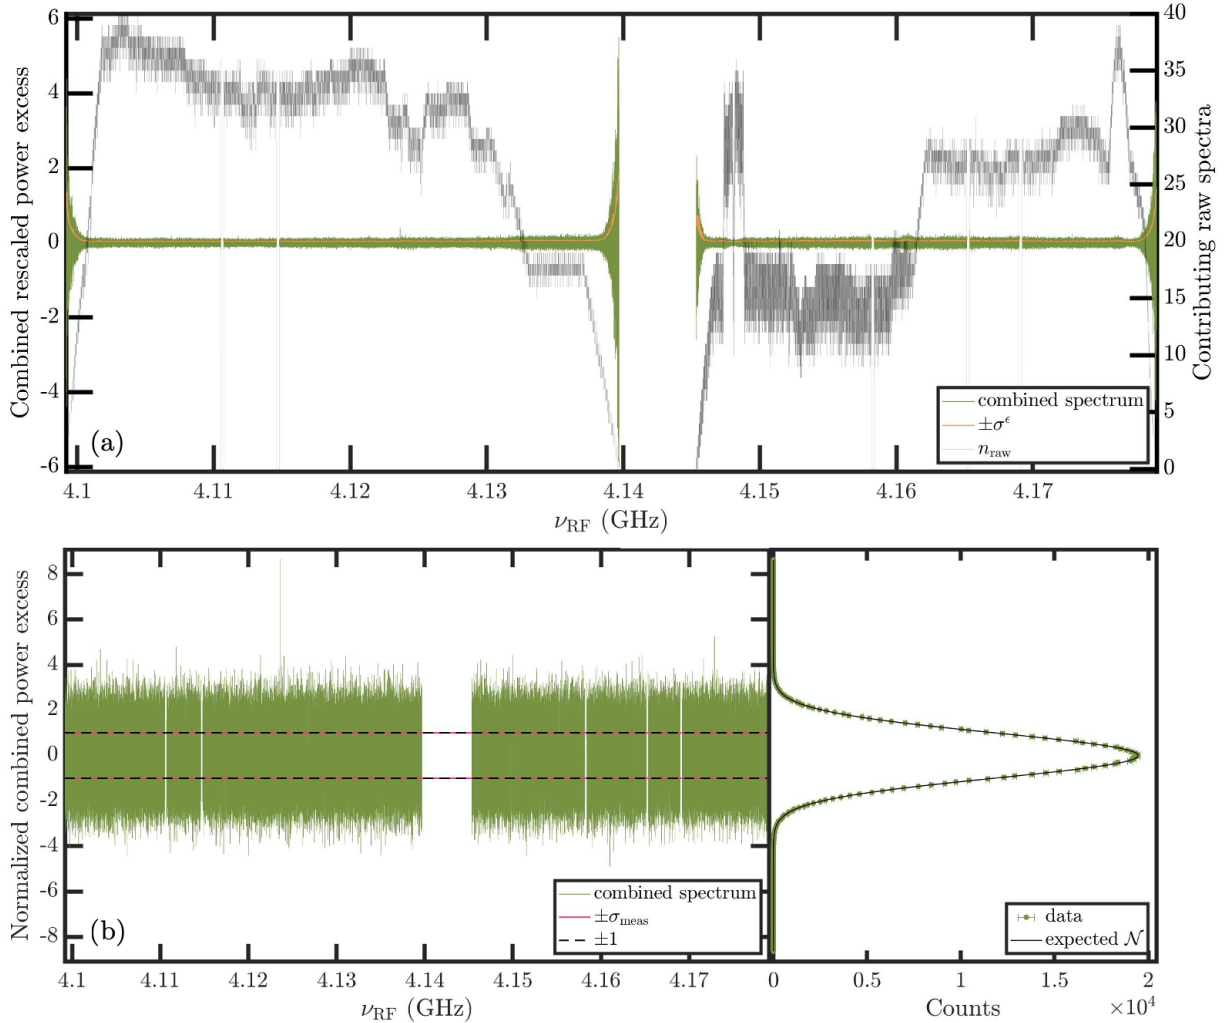


Figure 8.15: (a) HAYSTAC Phase II initial scan combined spectrum, obtained by summing the rescaled spectra with maximum likelihood (ML) weights such that a single-bin KSVZ axion has mean 1. The standard deviation σ^ϵ (orange line) therefore encodes sensitivity to the axion. Gray line/right axis shows the number of raw spectra that contribute to each combined spectrum bin, with higher axion sensitivity typically occurring where more spectra contribute. Small gaps in the initial scan combined spectrum were due to spurious, nonaxionic excesses which did not appear during rescans. The large gap beginning at 4.14 GHz is due to a transverse electric mode. (b) Alternate representation of the combined spectrum together with histogram of powers. The expected standard deviation σ^ϵ has been normalized out to give all bins unit standard deviation. The measured standard deviation σ_{meas} for the combined spectrum (pink line) matches this expectation (black, dashed line), and the data agrees well with the expected normal distribution.

8.2.5 Grand spectrum

To account for the axion lineshape, we apply ML weights to overlapping sets of combined spectrum bins as specified by Step 8 of Sec. 6.1.¹⁷⁴ The resulting (normalized, as in Fig. 8.15b for the combined spectrum) grand spectrum is shown in Fig. 8.16, where filter-induced correlations are corrected for according to Step 9 of Sec. 6.1. As in Fig. 8.15b, the axion sensitivity is encoded in the mean value μ^{KSVZ} that a KSVZ axion would impart to each bin, if its rest mass resided there. Here, we use the standard halo model lineshape of Eq. (2.14).




Figure 8.16: HAYSTAC Phase II initial scan grand spectrum, which accounts for the axion lineshape. The sensitivity to the axion is encoded by the mean μ^{KSVZ} (brown line) that a KSVZ axion with rest mass in each grand spectrum bin would have. Of particular initial interest was the spike far above μ^{KSVZ} at 4.1190395 GHz, which did not persist during rescanning. The standard deviation of the grand spectrum (pink line) matches the expected value (black, dashed line).

¹⁷⁴ The BPM framework makes it more natural to skip Step 7 of Sec. 6.1, and simply perform Step 8 on the combined spectrum. This is equivalent to setting $n_c \rightarrow 1$ and $n_r \rightarrow n_c n_r$ in Sec. 6.1, and enhances sensitivity to the axion very slightly while reducing complexity.

8.3 Results

The grand spectrum of Fig. 8.16 is combined with the rescan data (not shown) to report a constraint on the existence of axions within the scanned window. We will first discuss in Sec. 8.3.1 the results of the JILA analysis, using the BPM framework to examine the prior updates in the full two-dimensional space of frequency and coupling, the LU candidates, and the comparative advantage over a frequentist analysis. The official results, combining the JILA and Yale analyses, are presented and discussed in Sec. 8.3.2.

8.3.1 JILA analysis results




Figure 8.17: (a) Two-dimensional (axion frequency ν_a , axion-photon coupling g_γ) map of subaggregate prior updates U_s from the JILA analysis, with the 90% exclusion line ($U_s = 10\%$ contour) shown as the solid, blue line. The solid, red line shows the exclusion obtained using an frequentist threshold (FT) framework, with rescans marked as black dots. The bottom inset shows the top five least unlikely (LU) candidates with the ML couplings denoted by the corresponding orange circles in the main plot. None of the LU candidates posterior probabilities $P_{\text{LU}}^{(1)}$ occupy even 1% of the total prior probability \mathcal{P}_a of the window at their ML couplings. (b) The blue (red) solid line shows the aggregate prior update \mathcal{U} (exclusion \mathcal{E}) for the > 70 MHz region included within the two sets of dashed vertical lines spanned by dashed horizontal lines in (a). The dashed, blue horizontal lines in (a) indicate 90% exclusion for the Bayesian power-measured (BPM) framework, and the dashed, red line does so for the FT framework. The τ_{rel} axis (dashed, gray line at 90% exclusion) indicates the relative amount of scan time (+52%) it would have taken to achieve the BPM exclusion with the FT framework. The inset shows a closeup on the two frameworks' aggregate exclusions near 90% exclusion.

Applying the BPM framework of Ch. 7 to the HAYSTAC Phase II dataset produces the two-dimensional prior update map of Fig. 8.17a. Comparing to the Phase I reanalysis of Fig. 7.3 reveals a much cleaner result: there are no regions of very high subaggregated updates U_s , and the LU candidates, shown in the inset, never reach even $P_{\text{LU}}^{(1)}/\mathcal{P}_a = 1\%$ (compare to 20% in Fig. 7.3). The window quoted for aggregate exclusion in Fig. 8.17b is that spanned by the dashed horizontal lines denoting 90% aggregate exclusion in Fig. 7.3a. In both figures, blue lines indicate BPM results and red lines FT¹⁷⁵ results.

The aggregate exclusion properly accounts for the look-elsewhere effect with respect to exclusion for both Bayesian (Sec. 7.1.2) and frequentist (Sec. 6.3.2) analyses, and reveals a 52% more efficient scan using BPM. In particular, the frameworks are compared at the conventionally used 90% mark, where the scaling of relative scan time with coupling, Eq. (7.45), converts the couplings excluded at the 90% level over the whole range to relative scan times. The BPM framework’s advantage was aided by its operational flexibility (Sec. 7.3), which was used to rescan the 4.11190395 GHz initial scan excess (Fig. 8.16) multiple times.

The JILA results are not the official results of the Phase II HAYSTAC data run. Rather, they were averaged with the prior updates of the semi-independent Yale analysis. The two analyses agreed at better than the 1% level before averaging.

8.3.2 Official HAYSTAC Phase II results and conclusions

Averaging the JILA (Fig. 8.17) and Yale (not shown) analysis results, the final HAYSTAC Phase II results constrain the existence of axions with m_a between 16.96–17.12 and 17.14–17.28 $\mu\text{eV}/c^2$. The final results exclude axions with $g_\gamma \geq 1.38g_\gamma^{\text{KSVZ}}$. Figure 8.18a and b show the most important quantities plotted in Fig. 8.17. The results from the HAYSTAC quantum-enhanced data run are shown alongside other axion haloscope exclusion curves in Fig. 8.17c.

Quantum squeezing has allowed the HAYSTAC experiment to achieve a breakthrough in sensitivity by conducting the first sub-quantum-limited search for new fundamental particles. The

¹⁷⁵ The FT framework results shown here for comparison used the parameters of Ref. [106].




Figure 8.18: (a) Subaggregated prior updates U_s in the full two-dimensional parameter space of axion frequency ν_a and coupling g_γ achieved with the BPM analysis and shown in grayscale. The 10% prior update contour is shown in solid blue. The corresponding 90% aggregate exclusion level of $1.38g_\gamma^{\text{KSVZ}}$ is shown as a dashed blue line. (b) Aggregate prior update \mathcal{U} , Eq. (7.7), as a function of coupling over the entire window covered by the dashed blue line in (a). (c) Results of this work shown in blue alongside previous exclusion results from other haloscopes shown in red [89, 91, 93, 99, 102–105], orange [98, 101], purple [97, 100], and green [108]. Previous HAYSTAC results [106, 107] are shown in faded blue. The QCD axion model band [59] is shown in yellow, with the specific KSVZ and DFSZ model lines shown as black dashed lines. Figure adapted from Ref. [253].

SSR, capable of delivering 4.0 dB of noise variance reduction relative to vacuum with the large magnetic field turned on, has helped achieve world-leading sensitivity to axion dark matter in the $10 \mu\text{eV}/c^2$ mass decade. As in Ch. 5, the current achievable squeezing is limited primarily by microwave transmission losses between the two JPAs. As the development of superconducting quantum technology [218, 222] continues to improve microwave transmission efficiency, the scan

rate enhancement from squeezing will improve even without new experimental techniques [208], and advances in measurement techniques that remove nonreciprocal elements [252] offer a means of eliminating the majority of the transmission loss in the current SSR configuration.

8.4 Calibration measurements




Figure 8.19: Model used for HAYSTAC Phase II noise calibration measurements. The system is modeled for the purpose of calibration measurements with a cascaded series of circuit elements that perform discrete operations. A variable temperature stage (VTS) sources spectral density S_{in} . Boxes with Greek letters model loss, amplifiers (triangles) have the gains listed within them, and all added noise N_A is referred to the AMP. The cavity has resonance ν_c and coupling rates κ_m and κ_l , together giving rise to the susceptibility matrix elements $\chi_{ij}(\nu)$. The measured output spectral density is S_{out} .

Characterization of the sensitivity to the axion requires in-depth measurements of the system noise. To do this, we perform calibrations once every nine tuning steps. We use the model in Fig. 8.19 to calculate all quantities of interest. The model is chosen to balance level of detail against the number of independent measurements and reasonable assumption that can be made on the setup of Fig. 8.1. All spectral densities and gains in this section are single-quadrature unless otherwise specified.¹⁷⁶ The model consists of a variable temperature stage (VTS) modeled after that in Ch. 4 which sources Johnson-Nyquist noise with temperature-dependent spectral density S_{in} .

¹⁷⁶ The “1Q” subscripts of previous chapters where there was cause for ambiguity have here accordingly been dropped.

This noise is routed through the SQ (inferred squeezing $G_S \leq 1$)¹⁷⁷ to the cavity with susceptibility matrix elements (Sec. 5.1.1) χ_{ij} determined from its resonance ν_c , and measurement and (fictitious) loss port coupling rates κ_m and κ_l , respectively. From the cavity, the state goes through the AMP (gain G_A) and commercial amplifier chain (gain $G_{H'}$), led by a high-electron-mobility transistor (HEMT) amplifier. The system added noise spectral density N_A is referred to the input of the AMP, and any transmission loss that the state passes through between circuit elements is modelled with power transmissivities λ , ρ , and α as shown in the figure. The fluctuation-dissipation theorem [255] ensures that whatever fraction of the state is dissipated in these partially transmissive elements is replaced with Johnson-Nyquist noise at the fridge spectral density S_f . Single-quadrature spectral densities S (or N , which is used to represent added noise) have dimensions of photon number per unit bandwidth per unit time (i.e. dimensionless energy), corresponding to physical temperatures T at frequency ν via

$$S = \frac{1}{2} \left(\frac{1}{2} + \frac{1}{e^{h\nu/k_B T} - 1} \right) = \frac{1}{4} \coth \left(\frac{h\nu}{2k_B T} \right), \quad (8.3)$$

half the two-quadrature value. The Johnson-Nyquist noise spectral density of the cavity is $S_c \geq S_f$, to account for nonideal thermalization of the tuning rod.¹⁷⁸ Writing down the output spectral density S_{out} is therefore a simple matter of cascading the multiplicative and additive effects of all the elements:

$$S_{\text{out}} = G_{H'} G_A \left\{ N_A + (1 - \alpha) S_f + \alpha \left([1 - |\chi_{mm}|^2] S_c + |\chi_{mm}|^2 [(1 - \rho) S_f + \rho G_S (S_{\text{in}} \lambda + (1 - \lambda) S_f)] \right) \right\}. \quad (8.4)$$

This equation will enable the calibration calculations in Secs. 8.4.1 and 8.4.2.

¹⁷⁷ In this chapter, G_S is the “gain” along the squeezed quadrature, which is less than 1. It is ideally the inverse of $G_{S,1Q} \geq 1$ of previous chapters, which represented the single quadrature gain in the amplified quadrature of the SQ, but may be increased (i.e. worsened) due to distortion or added noise of the squeezer.

¹⁷⁸ The cavity spectral density cannot be determined from simply knowing the temperature of the tuning rod, as the tuning rod and cavity walls both participate in sourcing the “cavity” spectral density S_c . Therefore, though we measure the tuning rod to be at $T_{\text{rod}} = 225$ mK by slowly ramping up the fridge temperature and seeing where any excess PSD due to the rod vanishes, we must still independently measure S_c .

8.4.1 Squeezing and cavity spectral density measurements

The first step in the calibration protocol is a set of three power spectrum measurements (Fig. 8.20) to determine two experimental parameters: the spectral density S_c emerging from the cavity and the IF frequency (ν_{IF})-dependent squeezing $G_S(\nu_{\text{IF}})$ inferred over the analysis band. Measurement 1 is taken off cavity resonance with the squeezer turned off. Measurements 2 and 3 are taken on cavity resonance with the squeezer turned on and off, respectively. All three measurements are taken with $S_{\text{in}} = S_f$. From measurements 1 and 3 we calculate

$$S_c = \left\{ \left(\left[G_A^{(1)} (N_A + S_f) \frac{S_{\text{out}}^{(3)}}{S_{\text{out}}^{(1)}} \right] / G_A^{(3)} - N_A - (1 - \alpha) S_f \right) / \alpha - |\chi_{mm}|^2 S_f \right\} / (1 - |\chi_{mm}|^2), \quad (8.5)$$

from Eq. (8.4) with $G_S = 1$, where numerical superscripts denote quantities that are expected to vary between measurements 1–3. Other quantities are modeled as stable between measurements and the AMP gain is measured independently. Averaged over all operating points,

$$S_c = 0.41 \pm 0.02. \quad (8.6)$$

For a cavity at the base temperature of 61 mK, we would expect $S_c = S_f = 0.27$ from Eq. (8.3), close to the zero-temperature value of $1/4$. For a quantitative treatment of how the increased cavity spectral density affects the scan, see Appx. C. From measurements 2 and 3, with S_c determined, we calculate

$$G_S = \left\{ \left[\left(\left[G_A^{(1)} (N_A + S_f) \frac{S_{\text{out}}^{(2)}}{S_{\text{out}}^{(1)}} \right] / G_A^{(2)} - N_A - (1 - \alpha) S_f \right) / \alpha - [1 - |\chi_{mm}|^2] S_c \right] / |\chi_{mm}|^2 - (1 - \rho) S_f \right\} / (\rho S_f) \quad (8.7)$$

from Eq. (8.4). Using Eq. (8.7), we infer a typical¹⁷⁹

$$G_S \approx 0.1. \quad (8.8)$$

¹⁷⁹ In practice, states squeezed beyond the ideal deliverable squeezing set by the transmission loss do not differ greatly at the output. This is essentially why the scan rate enhancement in Fig. 5.2b rapidly levels out with respect to $G_{S,1Q}$. Fluctuations in the calibrated quantities, Eqs. (8.5), (8.7), and (8.9), tend to correlate or anti-correlate in such a way as to roughly preserve the total noise level.




Figure 8.20: (a) Sample calibration measurements taken at a representative operating point with (1, red) SQ off/off resonance, (2, green) SQ on/on resonance, and (3, blue) SQ on/off resonance. Sets of 300 adjacent $\Delta_b = 100$ Hz frequency bins are averaged together. Some of the large-scale structure owes to the spectral of the commercial amplification chain, and cancels out when ratioing these measurements to obtain parameters of interest. (b) Calculated values of the cavity spectra density S_c (in single-quadrature quanta) and squeezer gain $G_S \leq 1$ from the measurements in (a). The squeezer gain is treated as frequency-dependent. The low-frequency rise in squeezer gain may be due to an imperfect accounting of added noise by the added noise measurement, causing a slight increase in inferred squeezing.

8.4.2 Added noise measurement

The second calibration measurement is a hot-cold load measurement with SQ off ($G_S = 1$) similar to the protocol used in previous HAYSTAC results [90],¹⁸⁰ where the cold load is at the dilution refrigerator base temperature and the hot load is maintained at 333 mK and monitored by a Magnicon SQUID-based temperature sensor. Equation (8.4) gives us the IF frequency-dependent, single-quadrature added noise $N_A(\nu_{\text{IF}})$ referred to the input of the AMP as

$$N_A = \left\{ \left(\left[\left(\left[\lambda S_{\text{in}}^h + (1 - \lambda) S_f \right] \rho + (1 - \rho) S_f \right) |\tilde{\chi}_{mm}|^2 + (1 - |\tilde{\chi}_{mm}|^2) S_c \right] \alpha + (1 - \alpha) S_f \right) G_a^h - \frac{S_{\text{out}}^h}{S_{\text{out}}^c} \left([S_f |\tilde{\chi}_{mm}|^2 + (1 - |\tilde{\chi}_{mm}|^2) S_c] \alpha + (1 - \alpha) S_f \right) G_A^c \right\} / \left(G_A^c \frac{S_{\text{out}}^h}{S_{\text{out}}^c} - G_A^h \right), \quad (8.9)$$

¹⁸⁰ The Y-factor measurement of Ref. [90], in which $Y = S_{\text{out}}^h / S_{\text{out}}^c$, is a simplified version of the VTS noise calibration of Sec. 4.4.3, where the fridge temperature is not varied, and only two (one hot, one cold) VTS temperatures are used.

where $\lambda = -1.04$ dB is measured with a separate thermal sweep in which the fridge temperature is varied and $\alpha = \rho$ is assumed at half the independently measured value for the total intra-SSR loss of -2.0 dB.¹⁸¹ In Eq. (8.9), h and c superscripts denote quantities measured in the hot and cold load configurations, and the cold-load VTS input is simply S_f . The tilde on $|\tilde{\chi}_{mm}|^2$ denotes that we correct for the fact that, though this is an off-resonant measurement, the cavity's higher spectral density S_c is still close enough to be felt slightly, but only on one side of the LO.¹⁸² In Phase I of HAYSTAC [90], phase-insensitive JPA operation required the addition of at least a quarter of a quantum of noise per quadrature, $N_A \geq 1/4$. Now operating in phase-sensitive mode, the AMP adds an average

$$N_A = 0.03 \pm 0.02 \quad (8.10)$$

over the analysis band. This calibration measurement, together with those of Sec. 8.4.1, provides an accurate measurement of the total noise against which an axion signal must be measured.

In practice, due to the intrinsic limitations of any model with just a few discrete, cascaded elements, some of the calibration measurements for S_c , G_S , and N_A initially produce unphysical values. In order to still use these calibration data, we adopt a procedure whereby the AMP gains input to Eqs. (8.5), (8.7), and (8.9) are varied slightly relative to their measured values. When selecting a new value for the gains, we systematically bias towards values which produce lower scan rates: that is, we systematically bias our measurements in the conservative direction. An alternative approach to building in deliberate conservative bias in the face of imperfect calibration measurements would have been to unbalance our assignment of the total measured intra-SSR loss between ρ and α .¹⁸³

¹⁸¹ The loss measurement protocol is exactly that of Sec. 5.5.3. Vacuum is amplified first by the SQ with AMP off, and then by the AMP with SQ off at the same gain used for the SQ. The outputs for several different gains (to verify linearity) are compared at the output to determine how much of the state is lost between them. In Sec. 5.5.3, the intra-SSR loss was -1.6 dB. Here, the benefit from going to lower frequencies is outweighed by the disadvantages of a much less compact setup, constrained by the need for key circuit elements to be in the field-free region of the magnet, Figs. 8.4 and 8.5. Losses did not appear to increase when the magnet was ramped up.

¹⁸² Because Eqs. (8.5) and (8.9) are mutually dependent, the calibrations are in practice processed iteratively in order to converge on values for S_c , G_S , and N_A . The contribution of S_c to N_A is small, however, so a reasonable simplifying assumption for future analyses would be to set $S_c = S_f$ in Eq. (8.9), which would further eliminate the distinction between $|\tilde{\chi}_{mm}|^2$ and $|\chi_{mm}|^2$.

¹⁸³ Recall from Sec. 5.1.2 that, all else equal, loss in α is more harmful than loss in ρ , as in addition to degrading the squeezed state, it also attenuates the axion signal.

Chapter 9

Search for axion-like dark matter with the JILA eEDM experiment

Thus far, this thesis has focused on the QCD axion as a well-motivated resolution to the dark matter problem; however, as we saw in Ch. 2, QCD axions are far from the only particulate dark matter candidates, even among light particles. A popular and (as the name suggests) closely related class of candidates are axion-like particles (ALPs) [256]. ALPs, like axions, are athermally produced pseudo-goldstone bosons, but (by definition)¹⁸⁴ do not resolve the strong CP problem through the Peccei-Quinn mechanism of Sec. 2.3.2. For this reason, they are not as well motivated by particle physics as QCD axions, but they nonetheless could address problems such as baryogenesis, the cosmological constant, and small-scale structure formation [257]. Furthermore they are *not unallowed* by the laws of physics as we understand them, and the importance of solving the dark matter problem itself serves as motivation to look for any such particles within experimental reach. As with QCD axions, ALPs are expected to couple to normal matter, and may in fact do so more strongly than their namesakes. However, different models predict starkly different mass windows, with m_a plausibly ranging from 10^{-24} – 10^{-1} eV/ c^2 (10^{-10} – 10^{15} Hz) [257–260]. ALPs on the lower end of this mass range in particular could resolve problems raised by astrophysics pertaining to small scale structure [257,259,261,262]. Independent of mass, we will assume that ALPs, if present, have the same standard halo model (SHM) lineshape as we have used for QCD axions, Eq. (2.14).

One way to detect ALPs is through their effects on measurements of the electric dipole

¹⁸⁴ ALPs are somewhat loosely defined: they can alternatively be thought of as a super-class to QCD axions. This thesis adopts the perspective that “axion-like particle” and “QCD axion” are nonoverlapping categories. In this chapter, unlike in previous ones, the QCD modifier will be prepended to “axion” when there might otherwise be ambiguity.

moments (EDMs) of Standard Model particles.¹⁸⁵ An ALP of mass m_a may cause either a particle's EDM or its coupling to other particles to oscillate at frequency ν_a according to Eq. (2.9). The practical case for such ALP searches is made especially salient by the fact that the precision measurement of EDMs is a well-established field [263, 264], with an abundance of already existing, meticulously acquired, and potentially ALP-sensitive data [43, 265–270]. Inspired in large part by ALPs, physicists have therefore recently begun to take seriously the possibility of time-varying EDMs [271–275].

This chapter presents the work of Ref. [251], which uses the Bayesian power-measured (BPM) framework of Ch. 7 to perform one such search for ALP dark matter by considering its effect on the EDM of HfF^+ ions. The EDM measurements were originally carried out to look for a *static* electron EDM (eEDM) d_e in the JILA eEDM experiment [269, 276, 277]. A static eEDM above the Standard Model prediction of $|d_e| \leq 10^{-38} e \text{ cm}$ would constitute evidence of new physics beyond the Standard Model [278–282], but as of yet none has been observed in either the JILA eEDM experiment or the ACME experiment [270], currently the two most sensitive experimental platforms dedicated to finding a nonzero eEDM. But the absence of a static signal does not preclude the presence of an oscillating one, which could average to nearly zero over the long durations during which data is collected.

The goal is to see if an oscillating signal fits to the existing, time-stamped JILA eEDM data; the challenge is to properly identify and account for the myriad effects which shape an ALP signal in that context. After discussing how the JILA eEDM experiment may be sensitive to ALPs in Sec. 9.1, we will obtain the spectral content of its dataset in Sec. 9.2. Section 9.3 is the heart of this chapter: in it, we will fully quantify the no-ALP and ALP hypotheses, which entails accounting for not only the behavior of dark matter in an entirely different regime than that of haloscopes, but also for particulars of the experiment at hand. Finally, Sec. 9.4 presents the results of the full search, which spans frequencies over seven orders of magnitude, from 27 nHz to 400 mHz (10^{-22} –

¹⁸⁵ In Sec. 2.3.1, we already saw the direct role of one EDM — that of the neutron — to motivating the QCD axion. This chapter will look at measurements of a different EDM — that of the electron — to search for ALPs.

10^{-15} eV/ c^2 mass), and Sec. 9.5 discusses prospects for future searches. Observing no signature of an oscillating EDM in this range allows us to place the first laboratory constraint on the ALP-gluon coupling in the 10^{-17} – 10^{-15} eV/ c^2 mass range, as part of a total exclusion result spanning seven mass decades.¹⁸⁶ Moreover, the work presented here [251] is the first laboratory constraint to account for the stochastic nature of the ALP field.

9.1 ALP sensitivity of the JILA eEDM dataset

The JILA eEDM experiment is a Ramsey spectroscopy measurement [283] of the stretched-state splitting in both levels of a Stark doublet of $m_F = \pm 3/2$ states in $^{180}\text{Hf}^{19}\text{F}^+$ ions in an radio frequency quadrupole ion trap (or Paul trap) [284]. The ions are polarized in rotating magnetic and electric fields, and a spin precession measurement yields the Ramsey fringe frequency

$$f_R = 3g_F\mu_B B_{\text{rot}}/h + 3\alpha' f_{\text{rot}} \tilde{R}\tilde{B} - 2W_S C_S \tilde{D}\tilde{B} + 2\mathcal{E}_{\text{eff}} d_e \tilde{D}\tilde{B} + \dots, \quad (9.1)$$

which probes the $m_F = \pm 3/2$ energy splitting directly. In Eq. (9.1), g_F is the $F = 3/2$ state Landé g -factor, μ_B is the Bohr Magnetron, B_{rot} and f_{rot} are the strength and frequency of the rotating magnetic field bias, α' is a parameter describing Berry’s phase [285], $W_S = 20.4$ kHz is a molecule-specific structure constant, \mathcal{E}_{eff} is the effective electric field seen by the valence electron in HfF^+ , and C_S is scalar-pseudoscalar nucleon-electron coupling. Each measurement is performed in one of $2^3 = 8$ “switch states” determined by $\{\tilde{B} = \pm 1, \tilde{D} = \pm 1, \tilde{R} = \pm 1\}$, where \tilde{B} is the sign of the magnetic field, \tilde{D} is the populated Stark doublet, and \tilde{R} is the sense of the electric bias field rotation. A complete description of the experiment can be found in Ref. [269] and references therein.

Measurements of the eight switch states can be linearly combined in order to isolate different components of the measured frequency. Of particular interest from Eq. (9.1) is

$$f^{BD} = |2\mathcal{E}_{\text{eff}} d_e - 2W_S C_S|, \quad (9.2)$$

¹⁸⁶ For comparison, the Phase II HAYSTAC results of Ch. 8 spanned from 16.96–17.28 $\mu\text{eV}/c^2$, or 0.008 mass decades.

used to probe the eEDM itself. While the first term in Eq. (9.2) may therefore seem conspicuous, the known coupling mechanisms between d_e and ALPs are expected to be suppressed as m_a [275, 286, 287], deleterious to direct detection at the low masses considered in this chapter. Another candidate coupling is between an ALP field and the nuclear EDM, which would cause an oscillating molecular EDM, but this is also m_a -suppressed [288].

It is instead a putative ALP-gluon coupling interaction [289] that sources an m_a -independent effect on f^{BD} . In particular, the interaction is hypothesized to cause an oscillating nucleon-electron coupling C_S . Therefore, for the rest of this chapter, we will assume

$$d_e = 0, \tag{9.3}$$

consistent with the measurement of Ref. [269],¹⁸⁷ and reducing Eq. (9.2) to

$$f^{C_S} = 2W_S C_S. \tag{9.4}$$

Figure 9.1a shows the full time series, C_S -sensitive dataset of the JILA eEDM experiment, acquired in a $T \sim 1$ yr window centered in August 2016. The data is far from temporally uniform, having been taken in several discrete data runs culminating with the 11 day run in 2017 shown as the rightmost clump in Fig. 9.1a and in closeup in Fig. 9.1b. The data in this final window are the highest quality measurements, with typical variances of order one hundredfold lower than those of the other data points. For shorter timescale oscillations such as that illustrated in red in Fig. 9.1b, this obviates the need for using the remainder of the data in Fig. 9.1a. However, the full dataset is helpful in obtaining sensitivity to long timescale oscillators such as that shown in red in Fig. 9.1a. Figure 9.1c shows a histogram of the normalized f^{BD} measurements originally taken in Ref. [269]. They appear to follow a Gaussian distribution about the mean, which is itself consistent with zero (Footnote 187).

¹⁸⁷ The actual result reported in Ref. [269] is $d_e = (0.9 \pm 7.7_{\text{stat}} \pm 1.7_{\text{syst}}) \times 10^{-29}$ e cm (90% confidence), where the “stat” and “syst” subscripts denote statistical and systematic errors, respectively.




Figure 9.1: (a) Time-series data from the JILA electron electric dipole moment (eEDM) experiment, cast as measurements of $2W_S C_S$ versus time from August 16, 2016 (temporal center of the dataset). An example sine wave (red) is shown overlain with the data, illustrative of how the full dataset could in principle contain an oscillatory signal, even while averaging to zero. The full dataset shown in (a) is used for the low frequency analysis. (b) Closeup on the last clump of data in (a), collected over 11 days in 2017 and used for the high frequency analysis. A higher frequency sine wave (red) than in (a) is again imposed for illustrative purposes. (c) The full dataset in (a), cast as measurements of f^{BD} and histogrammed in normalized units relative to its mean. The data follows a Gaussian distribution (black curve), and Ref. [269] shows it to be consistent with $d_e = C_S = 0$. Panels (a) and (b) adapted from Ref. [251]. Panel (c) adapted from Ref. [269].

9.2 Spectral content of the JILA eEDM dataset

The nonuniform time sampling in Fig. 9.1a and b necessitates a method other than the Fourier transform for extracting the spectral content of the JILA eEDM dataset; least-squares spectral analysis (LSSA) [290, 291] will serve us in this role. LSSA minimizes the least-squares fit of a sine function of frequency $\nu = \omega/2\pi$ and unknown phase to a dataset. We use this to extract A_0 and B_0 as the values of the quadrature variables¹⁸⁸ \tilde{A} and \tilde{B} that uniquely minimize

$$\chi^2 = \sum_{i=1}^N \frac{1}{\sigma_i^2} \left[\tilde{A} \cos(\omega T_i) + \tilde{B} \sin(\omega T_i) - f^{C_S}(T_i) \right]^2. \quad (9.5)$$

In Eq. (9.5), T_i are the timestamps of the N data points $f^{C_S}(T_i)$ with variance σ_i^2 shown in Fig. 9.1a and b.

¹⁸⁸ This chapter will use several variants of the symbols A and B to represent closely related quantities: \tilde{A} and \tilde{B} are the variables as used for fits in Eqs. (9.5) and (9.8), A and B are the results of simulations to determine LSSA results, A' and B' are the rotated results of those simulations (Sec. 9.3.1), A_0 and B_0 are the actual, *measured* values for A and B , and A'_0 and B'_0 are the measured values for A' and B' , obtained by rotating A_0 and B_0 .

The results of the LSSA performed on the JILA eEDM dataset are shown in Fig. 9.2. The left portion of the plot comes from Eq. (9.5), and the right portion will be discussed below. The even spacing of sample frequencies ν makes for a higher density of data points along the horizontal axis (within left and right portions of the plot), hence the higher amplitudes $\sqrt{A_0^2 + B_0^2}$ on the right half of the plot are a simple manifestation of the look-elsewhere effect (Sec. 6.5). We can already observe that no oscillatory component stands out prominently from its peers. This is not necessarily an indication that there is no signal hidden within the noise, but it does mean that we should not expect to see any extremely obvious ALPs upon performing our BPM statistical analysis. The BPM analysis will enable us to account conservatively (Sec. 7.1.5) for what spikes are present, and put a constraint on ALPs over the seven frequency decades of Fig. 9.2.




Figure 9.2: Least-squares spectral analysis of the JILA eEDM dataset. The total magnitude of oscillation $\sqrt{A_0^2 + B_0^2}$ corresponds to the full time-series dataset of Fig. 9.1a and b. The low frequency analysis ($< 126 \mu\text{Hz}$) is obtained from Eq. (9.5), and the high frequency analysis ($> 126 \mu\text{Hz}$) from Eq. (9.11). Figure reproduced from Ref. [251].

In order to examine those seven frequency decades, the JILA eEDM dataset is broken up into two analysis regions: low frequency (27 nHz–126 μHz) and high frequency (126 μHz –400 mHz).¹⁸⁹ For the low frequency analysis, LSSA is performed directly on the f^{Cs} data. The high frequency analysis, conversely, involves breaking the individual data points in the 11 day 2017 data run into smaller measurements that compose them, in order to be sensitive to short timescale oscillations.

¹⁸⁹ The “high frequency” label is strictly relative to the lower frequencies considered within this chapter. It is roughly 10 orders of magnitude lower than the analyses of Chs. 7 and 8.

The Ramsey fringe frequency measurements f_R of Eq. (9.1) are composed of a set of asymmetry measurements

$$\mathcal{A}'_i = \pm \frac{N_i - \langle \text{Hf}^+ \rangle}{\langle \text{Hf}^+ \rangle}, \quad (9.6)$$

which are treated as individual data points in the high frequency analysis. In Eq. (9.6), N_i is a count of Hf^+ ions generated by state-selective dissociation of the HfF^+ ions in a single shot of the experiment and $\langle \text{Hf}^+ \rangle$ is the mean Hf^+ count for the set of shots containing the individual shot in question. The $+(-)$ sign corresponds to reading out the same (opposite) state as that which was prepared. These asymmetry measurements, in the static case, already oscillate over the free-evolution time t_R ($\sim 0.5\text{--}0.7$ s).

The effect of an ALP is then to modulate that existing oscillation at ν_a . In practice, this means fitting the asymmetry measurements to

$$\mathcal{A}^{\text{fit}}(t_R) = -C e^{-\gamma t_R} \cos(2\pi f_R t_R + \phi) + O, \quad (9.7)$$

where C is the fringe contrast, γ is the coherence decay rate, ϕ is the initial phase, and O is the offset. We account for a potential ALP-induced time (t) dependence as a small perturbation term

$$\delta f(t) = f^{C_S}(t) = \tilde{A} \cos(\omega t) + \tilde{B} \sin(\omega t) \quad (9.8)$$

in

$$f_R(t) = f_{R0} + M' f^{C_S}(t), \quad (9.9)$$

where f_{R0} is Eq. (9.1) with $d_e = C_S = 0$ and M' are the matrix elements used to transform between switch and parity bases [269]. Expanding Eq. (9.7) to second order in δf yields

$$\begin{aligned} \mathcal{A}^{\text{osc}}(t) &= -C e^{-\gamma t_R} \cos(2\pi [f_{R0} + M' f^{C_S}(t)] t_R + \phi) + O \\ &\sim -C e^{-\gamma t_R} \cos[2\pi f_{R0} t_R + \phi] + O \\ &\quad + C e^{-\gamma t_R} [2\pi t_R M' f^{C_S}(t)] \sin(2\pi f_{R0} t_R + \phi) + C e^{-\gamma t_R} 2 [\pi t_R M' f^{C_S}(t)]^2 \cos(2\pi f_{R0} t_R + \phi). \end{aligned} \quad (9.10)$$

For the high frequency analysis, we perform LSSA on the deviation of the N measured asymmetries, Eq. (9.6), from the modulated fit asymmetries, Eq. (9.10), minimizing

$$\chi^2 = \sum_{i=1}^N \frac{1}{\sigma_i^2} [\mathcal{A}_i^{\text{osc}} - \mathcal{A}'_i]^2, \quad (9.11)$$

where σ_i^2 are here the variances of \mathcal{A}'_i at acquisition times T_i . The results, cast as f^{Cs} quadrature amplitudes A_0 and B_0 through Eq. (9.8), are plotted summed in quadrature in the right portion of Fig. 9.2.

9.3 Quantifying the hypotheses

With the best-fit values from the LSSA in hand for a dense (relative to the scale on which frequencies are independent, Sec. 7.1.4) set of frequencies, we now need to precisely quantify the ALP and no-ALP hypotheses analogously to the axion and no-axion hypotheses shown in Fig. 6.1 in order to analyze the data with the BPM framework of Ch. 7. As in Sec. 6.2, it is helpful to think of the ALP (there, QCD axion) as modifying the noise-only, or no-ALP distribution. Therefore, we will first determine the no-ALP distribution in Sec. 9.3.1, and then account for the presence of a hypothetical ALP in Sec. 9.3.2. Key stochastic properties of the ALP distribution are discussed in Secs. 9.3.3 and 9.3.4, and an additional correction specific to the dataset is accounted for in Sec. 9.3.5.

9.3.1 No-ALP hypothesis

The no-ALP distribution is obtained for the low-frequency analysis through simulations based off of the timestamps T_i and noise of the real data shown in Fig. 9.1a. Assuming each of the data points has mean zero [269], standard deviation σ_i , a simulated time series of data is generated and LSSA performed to determine values of A and B for frequencies evenly spaced up to 126 μHz according to Eq. (9.5). 1,000 such simulations are performed in order to acquire reliable statistics. The distributions of the A and B quadrature observables are assumed to be Gaussian (Fig. 9.1c), with mean zero and standard deviations σ_A and σ_B , which are determined separately for each

frequency. The high frequency analysis is treated differently: a shuffling algorithm detailed in the Supplemental Material to Ref. [251] is used, but the essential results are the same: distributions for A and B characterized by σ_A and σ_B at each frequency.




Figure 9.3: Phase sensitivity of the JILA eEDM dataset. Simulated standard deviations $\sigma_{A'}$ and $\sigma_{B'}$ for the rotated quadratures for all frequencies across both low and high frequency analyses. The lower (higher) standard deviation, in pink (blue) belongs by construction to the more (less) sensitive rotated quadrature B' (A'). The difference between quadratures tends to diminish with frequency, and small differences at high frequencies (where blue mostly resides behind pink) are dominated by statistical precision, and not real phase sensitivity. Figure courtesy of Tanya Roussy.

The fact that (in the low frequency analysis) σ_A and σ_B are allowed to differ accounts for any phase sensitivity in the joint sampling profile of T_i and σ_i . Phase sensitivity matters because we do not know the phase of the ALP signal, if one exists, and a fluctuation along a quadrature to which we are sensitive should be taken more seriously than a fluctuation along a quadrature to which we are not. As a simple example of how phase sensitivity could manifest, imagine a dataset like that of Fig. 9.1a which only had measurements taken exactly one month apart. A sine wave — the \tilde{B} term in Eq. (9.5) — of frequency $\nu = 1/(2 \text{ months})$ starting at the first measurement would be perfectly insensitive to the data, whereas a cosine wave (\tilde{A} term) would contain all of the measurement’s sensitivity to any possible ALP signal. Typically, this effect diminishes with frequency (though there is no reason to expect it to be perfectly monotonic for a real sampling profile), where an oscillating wave (red lines in Fig. 9.1a and b) with many periods is unlikely to

behave in coordination with our very aperiodic T_i .

The simulated quadratures A and B are rotated into A' and B' , the generalized quadratures with the highest and lowest standard deviations, corresponding to the lowest and highest ALP-sensitivities, respectively. The insensitive (A') and sensitive (B') quadrature standard deviations $\sigma_{A'}$ and $\sigma_{B'}$ are plotted in Fig. 9.3. As expected, the phase sensitivity matters most at low frequencies. At high frequencies, the apparent small discrepancy between quadratures appears (unsurprisingly) to be fully explained by the standard deviation owing to finite simulation statistics (limited by computational time constraints). Therefore, in the high frequency analysis, the quadrature standard deviations are averaged with one another and also smoothed in a 36-bin running average. Aside from making evident the phase sensitivity of the data profile, rotation into A' and B' allows us to treat the measured values A'_0 and B'_0 at each frequency as statistically independent: as with phase-sensitive amplification (Sec. 3.4.6) of a vacuum state along the amplified versus deamplified quadrature, knowing A'_0 is useless for predicting B'_0 , and vice versa. The same is not generally true for A_0 and B_0 .

The no-ALP distribution is therefore the zero-mean two-dimensional Gaussian probability function (PDF)

$$f_{x,y}(x, y; \sigma_x, \sigma_y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp \left[-\frac{1}{2} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) \right]. \quad (9.12)$$

with $\{x, y, \sigma_x, \sigma_y\} \rightarrow \{A', B', \sigma_{A'}, \sigma_{B'}\}$. The probability density will be evaluated for each measured pair $\{A'_0, B'_0\}$ and compared to those of the ALP distribution, determined in Sec. 9.3.2.

9.3.2 ALP hypothesis

To quantify the ALP distribution, we first determine the effect an ALP would have on an otherwise noiseless measurement, and then combine that effect with the no-ALP distribution of Sec. 9.3.1. We begin by anticipating that the only effect of an ALP will be to increase the variances

$\sigma_{A'}^2$ and $\sigma_{B'}^2$ by some amount $\zeta\sigma_X^2$ as

$$\sigma_{AX}^2 = \sigma_{A'}^2 + \zeta\sigma_X^2 \quad (9.13)$$

$$\sigma_{BX}^2 = \sigma_{B'}^2 + \zeta\sigma_X^2, \quad (9.14)$$

where the ALP distribution will therefore be Eq. (9.12) with $\{x, y, \sigma_x, \sigma_y\} \rightarrow \{A', B', \sigma_{AX}, \sigma_{BX}\}$. To justify this statement, first note that our lack of knowledge of the ALP phase implies that the ALP cannot displace the means of the distributions of the quadrature observables. Put another way, the ALP will displace the mean for any one random realization of its phase, but averaged over all phases the effect is simply to broaden the two-dimensional no-ALP distribution. Furthermore, we know the distribution will remain Gaussian because the effect of the ALP on the quadratures is built of many small contributions from individual, interfering parts of the ALP field: a point derived in Ref. [75] which we will return to in Sec. 9.3.3. Finally, the ALP and the noise in the JILA eEDM experiment are uncorrelated¹⁹⁰ so their variances add straightforwardly as in Eqs. (9.13) and (9.14). Lastly, as a further consequence of not knowing the ALP phase, σ_X^2 must be the same for the two quadratures, even if the system noise, and thus σ_{AX}^2 and σ_{BX}^2 are not.

The construction of the ALP distribution is represented in Fig. 9.4 for the procedure discussed thus far: the unrotated quadratures A and B (Fig. 9.4a) are rotated to A' and B' to obtain the no-ALP distribution (Fig. 9.4b) before having their variances increased to account for the ALP (Fig. 9.4c). Having made and justified our ansatz for the form of the ALP distribution, we now set out to derive σ_X^2 , the variance of the ALP-only (i.e. noiseless) distribution, which we will arrive at in Eq. (9.24). Subsequently, we account for the attenuation factor ζ from Eqs. (9.13) and (9.14) in Secs. 9.3.4 and 9.3.5.

¹⁹⁰ Strictly speaking, while an ALP would not know about the noise in the JILA eEDM experiment, the reverse is not true. An ALP, if present, could have had the effect of very modestly increasing the standard deviations of the data points in Fig. 9.1. However, the effect is negligible compared to the signal such an ALP would have imparted.




Figure 9.4: (a) Two-dimensional Gaussian noise distribution at a frequency low enough to observe phase-sensitive effects (ellipticity) along the unrotated quadratures A and B . (b) No-axion-like particle (no-ALP) distribution, obtained from rotating the distribution in (a) into the independent, insensitive (A') and sensitive (B') quadratures. (c) ALP distribution obtained by adding the same variance to the no-ALP quadrature variances, reducing the ellipticity. The red dot marks a single point in quadrature space as it is manipulated from (a)–(c). Figure courtesy of Tanya Roussy.

9.3.3 Coherent ALP

We first contextualize and consider the effect of a coherent, oscillating ALP signal. Many of the ALP frequencies we wish to interrogate imply coherence times τ_c much longer than the duration of the measurement T (~ 1 yr for the low frequency analysis, ~ 11 days for high frequency): $\tau_c \gg T$. Recall from Sec. 2.4.3 that $Q_a \sim 10^6$, which remains true for ALP dark matter. Therefore, we will be in the regime of coherent ALPs so long an oscillating signal at the frequency of interest undergoes much fewer than 10^6 oscillations during the measurement. The axion haloscopes discussed elsewhere in this thesis are 10–17 orders of magnitude higher in frequency than the ALPs discussed in this chapter, with total measurement duration at each frequency typically lasting ~ 30 hrs. They are therefore in the opposite regime, $\tau_c \ll T$. The results reported on in this chapter [251] represent the first new laboratory constraint in the $\tau_c \gg T$ regime to properly account for ALP behavior.¹⁹¹ In Sec. 9.3.4, I will show results [251] that occupy the intermediate regime $\tau_c \sim T$, not directly addressed elsewhere.

If the ALP is coherent, then the displacement amplitude it imparts along the quadratures is

¹⁹¹ Reference [292] offers a rough correction to previous constraints of other publications based on the authors' findings (Footnote 192), but not a new constraint.

fixed for a given amplitude of its field; however, as already alluded to in Sec. 9.3.2, the amplitude of the ALP field is itself a stochastic quantity [292].¹⁹² Reference [75] provides an elegant derivation of the ALP field amplitude, which can readily be translated into the quadrature picture.¹⁹³ The derivation, which I reproduce here, begins with an individual axion-like “particle”¹⁹⁴

$$a_i(v_i, t) = \frac{\sqrt{2\rho_A/N_a}}{m_a} \cos \left[m_a \left(1 + \frac{v_i^2}{2} \right) t + \phi_i \right], \quad (9.15)$$

where ρ_A is the local ALP dark matter density¹⁹⁵ (Sec. 2.4.2), N_a is the contributing number of such particles (indexed by i), and v_i is the speed of each one. The phases ϕ_i for each particle are totally unknown: that is, distributed uniformly over $[0, 2\pi)$. The particles are considered sufficiently weakly interacting that we need not consider correlations between them.

In order to build up the local ALP dark matter field from the individual particles, Ref. [75] next partitions the particles into subsets Ω_j with practically indistinguishable speeds between v_j and $v_j + \Delta v$. There are amply many ALPs that there can simultaneously be a very large number of subsets, with each still treated classically. The expression for the coherent combination of all N_a^j particles in subset Ω_j is then

$$a_j(v_j, t) = \sum_{i \in \Omega_j} \frac{\sqrt{2\rho_A/N_a}}{m_a} \cos \left[m_a \left(1 + \frac{v_j^2}{2} \right) t + \phi_i \right]. \quad (9.16)$$

The only i -dependent term in the sum is ϕ_i , so the sum can be performed by recasting cosine as the real part of a complex exponential, and noting that

$$\sum_{i \in \Omega_j} e^{i\phi_i} = \alpha_j \exp(i\phi_j) \quad (9.17)$$

is the result of a two-dimensional random walk in the complex plane, with α_j describing the Rayleigh-distributed distance traversed in N_a^j steps. The phase $\phi_j \in [0, 2\pi)$ (as for ϕ_i) as there is

¹⁹² This fact has until recently been underappreciated in the ALP-exclusion literature. As a result, Ref. [292] points out that at least seven published constraints [273, 293–297] are overly aggressive by roughly an order of magnitude.

¹⁹³ Reference [75] describes the analysis framework for ABRACADABRA [110–112], which occupies the same $\tau_c \ll T$ regime as haloscopes, so the useful description of the coherent ALP therein is never directly used in the related work.

¹⁹⁴ “Particle” here means something more akin to “building block,” and quite distinct from a quantum of the ALP field. The ALP is still very much a wave in this derivation. More precisely, the particle-as-building-block is a collection of (actual) particles of sufficiently high occupancy as to behave classically, but sufficiently low occupancy as to make up an minute fraction of the total number of ALPs in a de Broglie volume.

¹⁹⁵ In this chapter only, we treat the local dark matter density as 0.4 GeV/cm^3 , in keeping with other published ALP exclusion results (Footnote 15). We again assume the ALP dark matter density ρ_A saturates the local dark matter density: $\rho_A = \rho_{\text{DM}}$.

no preferred direction for the random walk. Redefining $\alpha_j \rightarrow \alpha_j \sqrt{N_a^j/2}$ yields

$$a_j(v_j, t) = \alpha_j \frac{\sqrt{\rho_A N_a^j / N_a}}{m_a} \cos \left[m_a \left(1 + \frac{v_j^2}{2} \right) t + \phi_j \right] \quad (9.18)$$

where $\alpha_j \geq 0$ now has PDF

$$f_{\alpha_j}(\alpha_j) = \alpha_j e^{-\alpha_j^2/2}. \quad (9.19)$$

The final expression for the ALP dark matter field is the sum of all subsets Ω_j :

$$a(t) = \frac{\sqrt{\rho_A}}{m_a} \sum_j \alpha_j \sqrt{f_v(v_j) \Delta v} \cos \left[m_a \left(1 + \frac{v_j^2}{2} \right) t + \phi_j \right], \quad (9.20)$$

and is of direct use to us in both the fully coherent ($\tau_c \gg T$) and intermediate ($\tau_c \sim T$, Sec. 9.3.4) ALP regimes treated in this chapter. In Eq. (9.20), $f_v(v_j)$ is the ALP velocity distribution for the j^{th} particle subset. Equation (9.20) therefore has three random variables within the sum. Fortunately, we know all of their probability distributions: α_j is distributed according to Eq. (9.19), $f_v(v_j)$ according to Eq. (2.13) for SHM ALPs, and ϕ_j is uniformly distributed as above.

For coherent ALPs ($\tau_c \gg T$), there is no sum in Eq. (9.20) (or rather j sums to 1), as there are too few ALP oscillations throughout the measurement for it to be sensitive to the finite ALP linewidth. This has two implications: first, we can safely set $v_j = 0$ and $f_v(v_j) \Delta v = 1$ in Eq. (9.20). Second, we can translate this expression into our two-dimensional Gaussian picture for the ALP distribution (Sec. 9.3.2). This is because the two remaining random variables α_j and ϕ_j describe a Rayleigh distributed amplitude with random phase. One way to define a Rayleigh random variable is as $\sqrt{X^2 + Y^2}$, where X and Y are independent and identically distributed mean-zero, standard deviation- σ Gaussian random variables. The Rayleigh PDF of Eq. (9.19) corresponds to $\sigma = 1$. We can therefore directly relate σ_X of Eqs. (9.13) and (9.14) to the dark matter density.

The ALP-only variance σ_X^2 cares both about the local ALP dark matter density ρ_A and the ALP-gluon coupling C_G/f_a (where the symmetry-breaking energy scale f_a is factored out). Writing the oscillatory C_S term in the Ramsey fringe as

$$f^{C_S}(t) = 2W_s \eta \frac{C_G}{f_a} a(t), \quad (9.21)$$

where $\eta \approx 0.11$ is a coefficient relating C_S induced in a Hf nucleus to the QCD θ angle¹⁹⁶ [289, 298] yields

$$A' = 2W_S\eta \frac{C_G}{f_a} \alpha \frac{\sqrt{\rho_A}}{m_a} \cos(\phi) \quad (9.22)$$

$$B' = -2W_S\eta \frac{C_G}{f_a} \alpha \frac{\sqrt{\rho_A}}{m_a} \sin(\phi) \quad (9.23)$$

for the quadrature amplitudes of a pure (no noise), coherent (only one j term, with subscript left off) ALP signal. Equations (9.22) and (9.23) correspond to mean-zero Gaussians with

$$\sigma_X^2 = \left(2W_S\eta \frac{C_G}{f_a} \frac{\sqrt{\rho_A}}{m_a} \right)^2. \quad (9.24)$$

We now have a simple expression relating physical parameters to the pure ALP distribution, absent any attenuating effects captured by ζ in Eqs. (9.13) and (9.14).

The fact that we treat A' and B' as normally distributed even in the presence of an ALP fully accounts for the stochastic nature of the ALP. As pointed out in Ref. [292], “an unlucky experimentalist could even have near-zero field amplitudes during the course of their measurement.” This fact is naturally handled by the Rayleigh amplitude σ_j which has standard deviation roughly equal to its mean; however, it is better handled by the two independent Gaussian-distributed quadrature amplitudes (Sec. 9.3.1), as they can very easily accommodate asymmetric phase sensitivity through their noise-only variances $\sigma_{A'}$ and $\sigma_{B'}$.¹⁹⁷ Therefore, we choose to work in the quadrature picture.

9.3.4 Partially coherent ALP

In the regime of intermediate ALP coherence ($\tau_c \sim T$), the variance σ_X^2 added by the ALP is attenuated by the interference of different velocity components v_j in Eq. (9.20) over the du-

¹⁹⁶ An ALP may couple to gluons through the QCD θ angle and yet still not solve the strong CP problem if the gluon coupling does not generate the largest contribution to the ALP mass.

¹⁹⁷ It is worth noting that the statement about unlucky experimentalists is in a sense misleading: the coherent interference of all the velocity eigenmodes has the effect of almost always giving the ALP a far greater amplitude than it would have, all else equal, in the incoherent regime ($\tau_c \ll T$) occupied by haloscopes and similar platforms. We will see this effect partially manifest in the intermediate regime ($\tau_c \sim T$) of the next section. It is only by comparing to the unphysical case where the ALP deterministically (as in the incoherent case) assumes an amplitude that it could only achieve in the coherent case, that past publications have managed to report overly aggressive exclusion results (Footnote 192). In other words, while the experimentalist could in principle be unlucky in the sense described by Ref. [292], the wider ALP linewidths that could undo this effect, would, at constant dark matter density and coupling to baryonic matter, only make things worse.

ration of the measurement. First we ask: will this affect our measurement at all? The answer depends on both the ALP frequency (and by extension its coherence time, using $Q_a \sim 10^6$) and the measurement duration T . In Table 9.1 I compute the phase slip

$$\Delta\phi = 360^\circ \times \frac{T\nu}{Q_a} \quad (9.25)$$

between the frequency component at ν and that at $\nu(1+1/Q_a)$ over the duration of the measurement for the low and high ends of the low and high frequency analyses. There are frequency components beyond $\nu(1+1/Q_a)$, and the times used are approximate, so the values $\Delta\phi$ just give a rough indication of whether we are in the coherent, intermediate, or incoherent regime. Evidently, the low ends of both analyses are safely in the coherent regime of Sec. 9.3.3, whereas the high end of particularly the high frequency analysis should experience a significant interference effect¹⁹⁸ between different velocity components of the ALP field. The high end of the low frequency analysis undergoes a negligible phase slip for $T \sim 1$ yr, but the effect may have to be accounted if, for example, reanalyzing with another data run c.2020 (and/or further in the future) incorporated. The final row provides the approximate values for HAYSTAC Phase II (Ch. 8), which is emphatically in the incoherent regime, for comparison.

The quantitative effect of interfering velocity components was obtained through simulations. Using Eq. (9.20), along with Eq. (2.13) for the SHM velocity distribution and Eq. (9.19) for Rayleigh amplitudes, and independently generating each ϕ_j , the ALP field was numerically constructed at frequencies spaced across the entire range. The field is “fit” with LSSA to get (upon performing hundreds of such simulations) distributions of A' and B' with (ALP-only) variances $\sigma_{X,A'}^2$ and $\sigma_{X,B'}^2$. We define the decoherence attenuations as

$$\zeta_{d,A} = \frac{\sigma_{X,A'}^2}{\sigma_{X,A',\text{narrow}}^2} \quad (9.26)$$

$$\zeta_{d,B} = \frac{\sigma_{X,B'}^2}{\sigma_{X,B',\text{narrow}}^2}, \quad (9.27)$$

¹⁹⁸ We informally referred to this effect as “phase jarbling” while performing the analysis, and encourage future analysis teams working in the intermediate coherence regime to do the same: we found reliably that while few people had come across the word “jarble” (actual definition: “to wet; to bemire”; not listed in some dictionaries), most everyone knew what we meant by it.

	ν	T	$\Delta\phi$ (degrees)	coherence regime
low freq analysis, lowest freq	27 nHz	~ 1 yr	3×10^{-4}	coherent
low freq analysis, highest freq	126 μ Hz	~ 1 yr	1.4	coherent*
high freq analysis, lowest freq	126 μ Hz	~ 11 days	0.04	coherent
high freq analysis, highest freq	400 mHz	~ 11 days	140	intermediate
HAYSTAC Phase II	4.14 GHz	~ 30 hrs typ.	1.6×10^{11}	incoherent

Table 9.1: ALP coherence for high and low frequency ALP analyses and HAYSTAC Phase II. The frequencies ν and approximate measurement durations T are listed for five different cases: the lowest and highest frequencies of the low and high frequency analysis, and the HAYSTAC Phase II search of Ch. 8. The phase slips $\Delta\phi$ between frequency elements roughly one ALP (or QCD axion) linewidth ($Q_a \sim 10^6$) apart, calculated according to Eq. (9.25), provide a proxy for the amount of interference expected between different velocity components used to build up the dark matter field, Eq. (9.20). The different cases are sorted into coherence regimes for organizational clarity. *Note that the high end of the low frequency analysis is verging on the intermediate coherence regime.




Figure 9.5: Effect of ALP field decoherence on the ALP-only distribution variance for the high frequency analysis. Dots with statistical error bars are decoherence-induced attenuations ζ_d from simulations of partially versus fully coherent ALP variances. The blue line is a fit to a second-order polynomial, whose values were used in performing the analysis. For the low frequency analysis (not shown) $\zeta = 1$ for all ν . Figure reproduced from Ref. [251].

where $\sigma_{X,A',\text{narrow}}^2$ and $\sigma_{X,B',\text{narrow}}^2$ are the resulting variances from identical simulations for an infinitely narrowband ALP ($Q_a \rightarrow \infty$). In practice,¹⁹⁹

$$\zeta_{d,A} \approx \zeta_{d,B} \quad (9.28)$$

¹⁹⁹ The symmetric attenuation is no surprise, but we considered the quadratures separately as a check on the intuition that there should be no meaningful difference.

so we define

$$\zeta_d = \frac{\zeta_{d,A} + \zeta_{d,B}}{2}. \quad (9.29)$$

The simulations confirm that decoherence is negligible ($\zeta_d = 1$) throughout the low frequency analysis. The distributions for A' and B' remain Gaussian, and the simulated attenuation of the ALP-only distribution is plotted as a function of frequency in Fig. 9.5. The values used for the analysis come from a second-order polynomial fit. The effect only becomes significant for $\nu \gtrsim 100$ mHz.

9.3.5 Finite sample time correction




Figure 9.6: Effect of finite sample time on the ALP-only distribution variance. Dots with statistical error bars are finite sample time-induced attenuations ζ_s from simulations of instantaneous sampling versus real sampling durations. Blue lines are separate fits to the low and high frequency data, whose values were used in the analysis. Figure reproduced from Ref. [251].

The total attenuation ζ of the ALP variance in Eqs. (9.13) and (9.14) receives a second contribution from sampling over time periods t_R comparable to an ALP oscillation period. The effect is an averaging over a fraction of a period of the oscillation, causing a slight attenuation in sensitivity at the high end of both low and high frequency analyses. A procedure analogous to that used in Sec. 9.3.4 is sufficient to quantify the effect. In particular, we compare LSSA results on simulated sinusoidal data sampled much more finely than the timestamps in our real data set, and

on that same data averaged over the actual intervals of data collection (longer in the low frequency analysis). This test, averaged over all signal phases, is performed at several frequencies with the attenuations quantified for each rotated quadrature as

$$\zeta_{s,A} = \frac{\sigma_{X,A',\text{avg}}^2}{\sigma_{X,A'}^2} \quad (9.30)$$

$$\zeta_{s,B} = \frac{\sigma_{X,B',\text{avg}}^2}{\sigma_{X,B'}^2}, \quad (9.31)$$

where $\sigma_{X,A'}$ and $\sigma_{X,B'}$ ($\sigma_{X,A',\text{avg}}$ and $\sigma_{X,B',\text{avg}}$) are the variances obtained without (with) averaging. Results indicate

$$\zeta_{s,A} \approx \zeta_{s,B}, \quad (9.32)$$

prompting

$$\zeta_s = \frac{\zeta_{s,A} + \zeta_{s,B}}{2}, \quad (9.33)$$

analogous to Eq. (9.29). Figure 9.6 displays the results of these simulations, with separate fits used for the low and high frequency analyses. The total attenuation

$$\zeta = \zeta_d \zeta_s \quad (9.34)$$

is given as the product of the individual attenuations from the two effects, Eqs. (9.29) and (9.33), neglecting any small interaction the two might have with one another. This is the term that goes into Eqs. (9.13) and (9.14) to fully specify the ALP distribution.

9.4 ALP exclusion results

To obtain final results from the BPM framework, we simply calculate the total prior updates²⁰⁰ as

$$U_i \left(\nu_i, \frac{C_G}{f_a m_a} \right) = \frac{f_{x,y} [A'_0(\nu_i), B'_0(\nu_i); \sigma_{A'}(\nu_i), \sigma_{B'}(\nu_i)]}{f_{x,y} \left[A'_0(\nu_i), B'_0(\nu_i); \sigma_{AX} \left(\nu_i, \frac{C_G}{f_a m_a} \right), \sigma_{BX} \left(\nu_i, \frac{C_G}{f_a m_a} \right) \right]}, \quad (9.35)$$

²⁰⁰ As in Sec. 7.1.1, these reduce to the Bayes factors for our ALP dark matter. Unlike in Sec. 7.1.1, there is only one scan, so we skip the individual scan updates u_i and produce the total updates U_i .




Figure 9.7: (a) Logarithmically subaggregated Bayesian power-measured analysis results for the JILA eEDM data. Grayscale shows subaggregated prior updates U_s in the full two-dimensional space of ALP frequency ν_a and ALP-gluon coupling $C_G/(f_a m_a c^2)$, with darker (lighter) colors indicating increased (decreased) belief in the ALP. Strings of blue dots show subaggregated 95% exclusion for each decade, and the dashed, blue line does the same for the entire frequency window. (b) Aggregate prior update \mathcal{U} versus coupling for the entire seven-decade frequency window. A small ($\mathcal{U} \approx 2.5$) peak likely due to technical noise is not enough to indicate the presence of an ALP. The dashed, blue line corresponding to that in (a) marks $\mathcal{U} = 5\%$ (95% aggregate exclusion). The dashed, black line shows the coupling at $\mathcal{U} = 5\%$ occurs. Logarithmically uniform priors are used throughout (a) and (b). Figure adapted from Ref. [251].

where dependencies on frequency and coupling have been made explicit in all terms, and the two-dimensional Gaussian PDF is given in Eq.(9.12). Note that the coupling to ALPs only enters through the ALP distribution variances. The total updates are then aggregated, Eq.(7.7), here using

$$\mathcal{U}\left(\nu_i, \frac{C_G}{f_a m_a}\right) = \frac{\sum_i U_i \epsilon/\nu_i}{\sum_i \epsilon/\nu_i}, \quad (9.36)$$

where ϵ/ν_i is our logarithmically uniform prior belief in the ALP, Eq.(7.14). The U_i are also subaggregated (U_s) using Eq.(9.36) into the 100 logarithmically spaced frequency bins shown in Fig.9.7a. The corresponding aggregated results \mathcal{U} of Fig.9.7b actually go upwards to $\mathcal{U} \approx 2.5$ at intermediate couplings before plummeting downwards. As is apparent in Fig.9.7a, the upward spike in \mathcal{U} comes from contributions at many frequencies. Even if one of these happens to be an ALP (unlikely, given at most a 2.5-fold increase in our very low priors), the others are indicative of system noise at sensitivities we would like to probe. Since the primary science goal of the JILA eEDM

experiment is (as the name suggests) to detect a nonzero eEDM, such features (and their successors in future datasets) may have to be tolerated. It is therefore important to keep in mind, per the discussion in Sec. 7.1.5, that the prior updates should really be viewed as conservative with respect to exclusion. Therefore, the final result, 95% BPM exclusion for $C_G/(f_a m_a) \geq 4.79 \times 10^{15} \text{GeV}^{-2}$ from 10^{-22} – 10^{-15} eV, can be safely believed.

The exclusion of Fig. 9.7 is plotted alongside other laboratory and astrophysical constraints in Fig. 9.8. The results of this chapter, shown in pink, include the first laboratory constraints on the ALP-gluon coupling from 10^{-17} – 10^{-15} eV/ c^2 .




Figure 9.8: Astrophysical and laboratory limits on the ALP-gluon interaction. The laboratory constraints of this chapter (pink region) are shown alongside other constraints at 95% confidence level. The purple region corresponds to ALPs with de Broglie wavelengths smaller than dwarf galaxies [260]. The blue region is from laboratory results on neutron EDM (nEDM) measurements [273]. The green region represents constraints from big bang nucleosynthesis [299]. The yellow region represents constraints from supernova energy loss bounds [249, 300]. Figure adapted from Ref. [251].

9.5 Prospects for future JILA eEDM ALP searches

Future ALP searches will be straightforward to perform with data from the next generation eEDM experiment, expected to be an order of magnitude more sensitive; however, two items are worth cataloguing at present for when those data arrive. First, the time duration T over which data was taken will extend several years. Unlike for haloscope results, where prior updates should simply

be multiplied between separate data runs, update-multiplication will not be the information-optimal way to integrate new data at low frequencies. Rather, where the ALP field exhibits coherent effects over the duration of data acquisition, a “very low frequency” analysis should be performed on the entire combined dataset.²⁰¹ At higher frequencies, updates may simply be multiplied, as there would be no benefit, and significant computational cost, from a combined reanalysis. In the intermediate case, as long as computational capabilities allow it, the coherent (“very low frequency”) approach is superior. Second, the choice to stop at 27 nHz on the low end was not necessary. Though a coherent ALP with oscillation period much longer than T could have its sinusoid be at a zero-crossing, the argument is analogous to that of the stochastic ALP field being at a null (Sec. 9.3.3): it does not prohibit us from accounting for the fact and conducting a search nonetheless. Acting on both of these items should, in conjunction with more sensitive data, produce more stringent constraints (or, preferably, a signal) at very low, low, and high frequencies. The techniques of this chapter may also straightforwardly be applied to other existing, ultralight ALP-sensitive datasets.

²⁰¹ In principle, the same could even be done between the datasets of multiple different ALP-sensitive experiments. However, our experience with the JILA eEDM dataset recommends caution: often specific details of the measurement (e.g. Sec. 9.3.5) are required in order to produce an accurate accounting.

Chapter 10

Conclusion and outlook



Figure 10.1: Axion namesake-inspired representation of the two Josephson parametric amplifiers (JPAs; one performing the inverse gain operation of the other) and axion cavity comprising the SSR-equipped haloscope.

The mysterious nature of dark matter remains among the largest and most prominent problems in all of modern scientific inquiry. The most popular hypotheses now point to a dark sector of abundant particles, which Occam’s razor suggests may in fact all be copies of the same particle. Among these, the QCD axion (Ch. 2) is well motivated, detectable, and as-yet not meaningfully disfavored by the results of direct detection efforts. The foremost goal of this thesis, and of the axion detection community more broadly, is to change that status quo: either by making a successful detection or by sweeping out the QCD band [59] accurately and thoroughly, and discovering its absence.

10.1 Conclusion

I have presented two separate but complimentary means for advancing towards that ambitious but worthwhile goal. The first, demonstrated in Ch. 5, relies on the Josephson parametric amplifiers (JPAs) of Ch. 3 and builds from the experimental techniques of Ch. 4 [204]. With it, I showed how two JPAs in a squeezed state receiver (SSR) configuration²⁰² [208] can prepare, deliver, and read out a squeezed state of the microwave field in order to circumvent the standard quantum limit on haloscope noise performance. Presently, the primary limitation on this method comes from experimental losses in the system, 1.6 dB in the demonstration setup of Ch. 5. The second improves upon the statistical analyses historically used for axion detection, which are discussed in Ch. 6. The Bayesian power-measured (BPM) framework of Ch. 7 makes full use of the information content of the processed haloscope measurement to test the axion hypothesis. It provides a more intuitive and flexible means of analyzing axion-sensitive data that what has traditionally been done. In Ch. 9 we saw its versatility in accommodating axion-like particle (ALP)-sensitive data as much as 17 orders of magnitude in mass below its original region of intended use [251], where the ALPs behave very differently than the QCD axions searched for by haloscopes.

The primary result of this thesis, that of Ch. 8, is the combined use of these two innovations to speed up a real haloscope search roughly threefold [253]. The application of squeezing in particular delivers on the long-held promise of quantum meteorology towards improving fundamental measurements [301], now in the microwave regime in addition to the optical [231]. The SSR-equipped haloscope of Ch. 8 was limited by 2.0 dB of transmission losses, similar to in Ch. 5, allowing for a nearly twofold speedup from squeezing alone.

The result of a quantum-enhanced axion dark matter search is of particular significance not simply for the speedup itself, nor for the relatively small scientific contribution to our knowledge of where in parameter space the axion is not (Fig. 8.18), but because it represents the first time a search for new particles has benefited from a quantum-enhanced measurement. Axion haloscopes

²⁰² For a representation of the SSR inspired by the dish detergent namesake of the axion [48], see Fig. 10.1.




Figure 10.2: (a) Approximate decrease of haloscope total noise levels as compared to the quantum limit over time. The figure shows the noise levels quoted in axion haloscope publications which marked significant advances in the field: the first haloscope search (RBF 1987 [97]), the first ADMX result (ADMX 1998 [89]), the first use of a SQUID-based amplifier in a haloscope (ADMX 2010 [104]), the first such use of a JPA (HAYSTAC 2017 [106]), and the work of Ch. 8 of this thesis (HAYSTAC 2020 [253]). We cannot reduce noise below zero (gray region), but there is a nearly unlimited amount to be won in the region between zero noise and the quantum limit (thin yellow region between the gray and white regions). (b) Approximate noise contribution to scan rate for the same publications marked in (a) and plotted along a logarithmic axis. Scan rate, Eq. (2.16) improves quadratically with noise relative to the quantum limit, making the linear noise decrease (1987–2017) of (a) a supralinear trend for the haloscope’s primary figure of merit. There is no bound on scan rate past the quantum limit (yellow region).

are among a small class of experimental platforms which had become limited primarily by quantum

noise fluctuations [91, 106, 229, 231]. Looking towards the future, it is unlikely that haloscopes (and long-baseline gravitational wave detectors along with them) will remain the exception and not the rule. Figure 10.2a illustrates this point starkly within the historical context of haloscope development. For the past three decades, axion noise levels (in quanta) have fallen almost perfectly linearly to the point where they were near-quantum-limited [106] prior to the work of this thesis. The corresponding effect on scan rate, shown in Fig. 10.2b, is by contrast quite nonlinear. If the improvement of haloscope noise levels, nearly the sole driver of scan rate enhancement (Sec. 2.5.3) since 1987,²⁰³ were to stop at the quantum limit, the time to sweep out the axion parameter space might well be prohibitive. The blue dots in Fig. 10.2 correspond to the work of this thesis, which enters into the yellow region,²⁰⁴ beyond the quantum limit. As indicated in Fig. 10.2b, there is no bound on scan rate in this new regime. Breaking through the quantum limit thus invites a new era of fundamental physics searches in which noise reduction techniques yield unbounded benefit rather than the diminishing returns of approaching the quantum limit. In Sec. 10.2, I discuss how axion physicists might most efficaciously accept that invitation.

10.2 Outlook

The two innovations in this thesis have distinct virtues. As the title of Ch. 7 suggests, the BPM framework is a terminal destination along the road to improving haloscope analyses, at least in terms of the axion-pertinent information content they can extract from haloscope data. However, its zero-cost applicability and tangible benefits recommend it strongly for rare event searches including especially axion haloscope experiments, and its improved flexibility may facilitate advances in scan strategy. Conversely, the SSR of Ch. 5 marks the beginning of an era of quantum enhanced searches

²⁰³ Notice that the left edge of Fig. 10.2 is 1983, the year the haloscope concept was first proposed by Pierre Sikivie [54]. It took only four years to perform the first haloscope search [97]. The right edge of Fig. 10.2 is 2024, four years after the writing of this thesis. In those four years, it is possible, but unlikely, that the axion may finally be discovered. Perhaps a stronger possibility is that the technology which will one day lead to a detection will be invented or advanced.

²⁰⁴ There are several ways one might perform the accounting in Fig. 10.2. I opt for a conservative one with respect to the central result of this thesis, in which the contribution of poor tuning-rod thermalization (Appx. C) adds back much of the noise which squeezing removes. Even with this conservative accounting, the noise remains under the quantum limit.

for new particle phenomena. While its delivered twofold scan rate enhancement (Ch. 8) is on its own an important step forward, as discussed above, its real advantage is in demonstrating the viability of sub-quantum-limited particle experiments.

Improving haloscope scan rates through improved scan strategy is a zero-cost advance enabled by and related to the operationally flexible BPM framework. Scan strategy has received a small amount of systematic treatment in the haloscope literature [95,126], but remains largely unexplored. Meanwhile, the current protocol adopted by most haloscopes — perform an initial scan, rescan peaks above a given threshold and/or below a certain sensitivity for roughly long enough to duplicate the sensitivity of the initial scan, perform manual interrogation (Footnote 145) if some number of rescans pass threshold — is not optimized. The discussion in Sec. 7.1.3 illuminates roughly how one might use the prior update in a practical haloscope setting to decide when to transition from rescans to manual interrogation. In addition, a more optimal scan strategy should account for larger power excesses with (all else equal) longer rescans. A further consideration in determining rescans could be the (cavity sensitivity profile-weighted) posterior belief that the axion is in *any* frequency bin within a cavity linewidth. Real-time data processing and analysis could significantly improve live times by informing the haloscope to simply wait longer near frequencies exhibiting a power excess, rather than the time-expensive step of going back and rescanning them. These and related considerations recommend constructing a value function based on posterior belief to determine initial scan, rescan, and manual interrogation strategy. A value function-based approach to scan strategy, enabled by the inter- and intrascan freedoms (Sec. 7.3) of the BPM framework, would also need as inputs realistic time cost estimates for various haloscope operations, which would differ between experiments. If implemented in software, it could passively improve scan rates for existing haloscopes and other narrowband, tunable detectors.

The historical trend shown in Fig. 10.2 suggests that the strongest prospect for enhancing haloscope scan rates is through continued improvement to noise performance. The SSR concept in this paper suggests two avenues for doing this.²⁰⁵

²⁰⁵ These are not the only approaches available. At higher frequencies, for example, it would likely be advantageous

The first way to improve upon the SSR is by simply waiting. While such a suggestion may sound frivolous, it is in fact one of the strongest arguments for using quantum measurement to reduce noise, rather than, for example, developing multiplexed cavity readout technologies to improve signal [302]. Put simply, there are very few technologists outside of the axion haloscope community working on coordinating the tuning and efficient readout of TM_{0n0} -like modes of multiple cryogenic cavities immersed in large magnetic fields. Conversely, the technologies that are used to improve noise and loss performance are the same ones that are constantly being developed and improved by the superconducting qubit community in order to achieve ever more impressive results [303, 304]. As is shown in Sec. 5.5.3 and Fig. 5.2, the marginal benefit to scan rate from each 0.1 dB of reduced transmission loss grows nonlinearly in the SSR setup. Specific technologies that could expedite this progress include on-chip microwave circulators [222], wireless parametric amplifiers [305], and directional amplifiers [218, 306].

An alternative approach is to get rid of any nonreciprocal elements in the signal path altogether. Reference [252] explores one way of doing this which is the natural dual to the SSR approach, and has one considerable advantage over it. Rather than *deamplifying* (or squeezing) the noise reflected off of the cavity, which creates a fragile quantum state which must be transported around the haloscope setup, the alternative is to *amplify* the noise coming out of the cavity, and with it an axion signal. The concept is illustrated in Fig. 10.3. Figure 10.3a shows the noise power spectral densities (PSDs) along a single quadrature in a standard haloscope configuration. The cavity and reflected noise PSDs complement one another, and, neglecting any added noise or thermalization problems (Appx. C), the total noise spectral density is flat. The SSR reduces the reflected noise, leaving the cavity noise untouched (Fig. 10.3b). In an idealized picture, it is equivalent to amplifying the cavity noise relative to the reflected noise (Fig. 10.3c). In practice, the latter approach is far more robust to transmission losses, as the axion-imprinted state leaves the apparatus itself amplified, along with the noise fluctuations which also originated within the cavity.

In practice, this approach can be implemented with a setup like that of Fig. 10.3d. The axion

to measure in the photon basis [83, 84, 163].

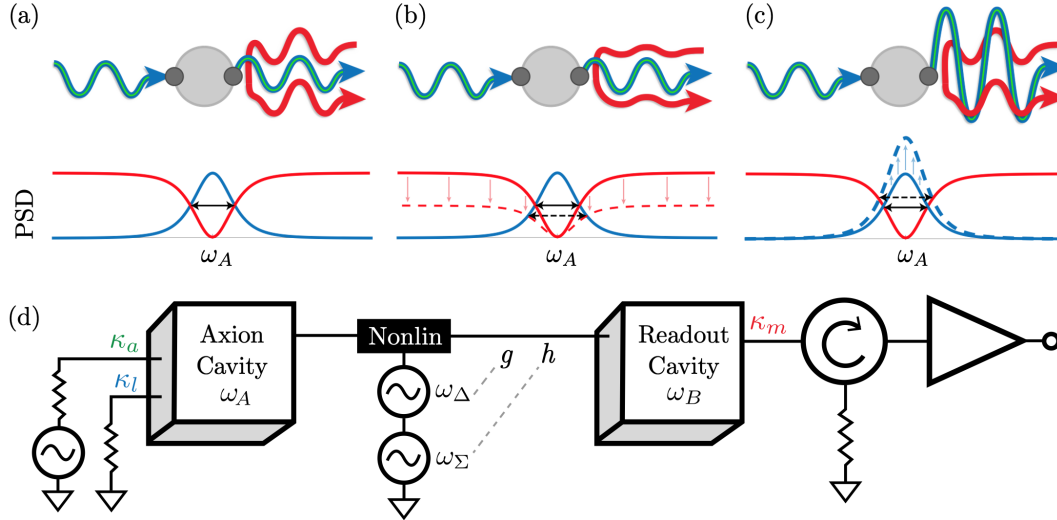


Figure 10.3: (a)–(c) Manipulations of vacuum noise within a single quadrature. (a) shows the noise profile of a standard haloscope. The reflected Johnson-Nyquist noise (red oscillation) is the same size as the Johnson-Nyquist noise coming from within the cavity (blue oscillation). Since a feeble axion signal (green oscillation) comes in fixed ratio with the cavity noise, the goal is to maximize the ratio of cavity-to-reflected power spectral densities (PSDs), illustrated below with the cavity's Lorentzian transmission and reflection profiles about the resonance frequency ω_A of its mode (gray circle). The squeezed state receiver squeezed reflected noise without affecting cavity noise as shown in (b) (compare to Fig. 8.9). An alternative, shown in (c), is to amplify cavity noise relative to reflected noise, which has the advantage of making the resulting output state more robust to transmission losses. (d) Schematic of a physical implementation of the concept in (c). The axion cavity with loss rate κ_l and axion-coupling rate κ_a is coupled through a nonlinear element (Nonlin) to a readout cavity with resonance $\omega_B \neq \omega_A$. The axion is modeled as a fictitious microwave generator as in Sec. 5.3. By modulating at the sum and difference frequencies ω_Σ and ω_Δ , respectively, simultaneous two-mode entanglement (rate h) and state swapping (rate g) interactions are generated between the cavity modes. The result is an amplification of the cavity noise emerging from the readout cavity (external coupling rate κ_m), *before* it encounters a nonreciprocal element and subsequent amplification. Figure adapted from Ref. [252].

cavity at ω_A is coupled to a second, readout cavity at ω_B . The states of the two are entangled at rate h by pumping a nonlinear element (such as SQUID array) between them at their sum frequency

$$\omega_\Sigma = \omega_A + \omega_B, \quad (10.1)$$

while a difference frequency

$$\omega_\Delta = |\omega_A - \omega_B| \quad (10.2)$$

pump simultaneously swaps the cavity states at rate g [307]. When the rates are equal, the system

behaves as a phase-sensitive amplifier [308, 309]. Figure 10.3d models coupling rates κ_a and κ_l for the fictitious axion and loss ports, and κ_m for the measurement port, analogous to Fig. 5.1. The state decaying out of the measurement port encounters a circulator (followed by a JPA or other leading amplifier) only after it has been amplified. This is essential, as the circulator has the effect of routing reflected signals from imperfections in the microwave elements before and after it away. Conversely, an imperfection between the two cavities just changes their coupling slightly, which can be compensated for via their antenna couplings. The reader is referred to Ref. [252] for a full input-output theory treatment of the system, alongside practical considerations for near-term implementation.²⁰⁶ As a final note, this approach can itself be complementary to the SSR: there is no reason one cannot both reduce off-resonant noise and increase on-resonant noise (and with it, signal), all at the same time.

For now, the axion, if it exists, remains well cloaked in its vast parameter space. Quantum measurement techniques, together with better classical information handling, has the potential to change that.

²⁰⁶ We call this approach the Cavity Entanglement And Swapping Experiment For Improved Readout Efficiency, or CEASEFIRE.

Bibliography

- [1] W. T. Kelvin, Baltimore lectures on molecular dynamics and the wave theory of light. London, Baltimore: C.J. Clay and Sons; Publication agency of the Johns Hopkins University, 1904, p. 274. [Online]. Available: <https://catalog.hathitrust.org/Record/011929196>
- [2] Ade, P. A. R., Aghanim, N., Arnaud, M., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A. J., Barreiro, R. B., Bartlett, J. G., Bartolo, N. et al., “Planck 2015 results,” Astron. Astrophys., vol. 594, Sep 2016. [Online]. Available: <https://dx.doi.org/10.1051/0004-6361/201525830>
- [3] A. A. Penzias and R. W. Wilson, “A measurement of excess antenna temperature at 4080 Mc/s.” Astrophys. J., vol. 142, pp. 419–421, Jul 1965.
- [4] G. Bertone and D. Hooper, “History of dark matter,” Rev. Mod. Phys., vol. 90, no. 4, Oct 2018. [Online]. Available: <https://dx.doi.org/10.1103/RevModPhys.90.045002>
- [5] J. G. de Swart, G. Bertone, and J. van Dongen, “How dark matter came to matter,” Nat. Astron., vol. 1, no. 3, Mar 2017. [Online]. Available: <https://dx.doi.org/10.1038/s41550-017-0059>
- [6] F. Zwicky, “Die rotverschiebung von extragalaktischen nebeln,” Helv. Phys. Acta, vol. 6, pp. 110–127, Jan. 1933.
- [7] F. Zwicky, “On the masses of nebulae and of clusters of nebulae,” Astrophys. J., vol. 86, p. 217, Oct. 1937.
- [8] E. Hubble and M. L. Humason, “The velocity-distance relation among extra-galactic nebulae,” Astrophys. J., vol. 74, p. 43, Jul. 1931.
- [9] S. Chandrasekhar, “The time of relaxation of stellar systems. III.” Astrophys. J., vol. 93, p. 323, Mar. 1941.
- [10] S. Thornton and J. Marion, Classical Dynamics of Particles and Systems, 5th ed. Brooks/Cole, 2004, pp. 189–190. [Online]. Available: <https://books.google.com/books?id=HOqLQgAACAAJ>
- [11] V. C. Rubin and J. Ford, W. Kent, “Rotation of the Andromeda nebula from a spectroscopic survey of emission regions,” Astrophys. J., vol. 159, p. 379, Feb. 1970.
- [12] D. H. Rogstad and G. S. Shostak, “Gross properties of five Scd galaxies as determined from 21-CENTIMETER observations,” Astrophys. J., vol. 176, p. 315, Sep. 1972.

- [13] R. N. Whitehurst and M. S. Roberts, “High-velocity neutral hydrogen in the central region of the Andromeda galaxy,” *Astrophys. J.*, vol. 175, p. 347, Jul. 1972.
- [14] M. Bartelmann, “The dark universe,” *Rev. Mod. Phys.*, vol. 82, no. 1, pp. 331—382, Feb 2010. [Online]. Available: <https://dx.doi.org/10.1103/RevModPhys.82.331>
- [15] D. Samtleben, S. Staggs, and B. Winstein, “The cosmic microwave background for pedestrians: A review for particle and nuclear physicists,” *Annu. Rev. Nucl. Part. Sci.*, vol. 57, no. 1, pp. 245—283, Nov 2007. [Online]. Available: <https://dx.doi.org/10.1146/annurev.nucl.54.070103.181232>
- [16] M. Milgrom, “A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis.” *Astrophys. J.*, vol. 270, pp. 365–370, Jul. 1983.
- [17] M. Milgrom, “A modification of the Newtonian dynamics - Implications for galaxies.” *Astrophys. J.*, vol. 270, pp. 371–383, Jul. 1983.
- [18] D. Clowe, A. Gonzalez, and M. Markevitch, “Weak-lensing mass reconstruction of the interacting cluster 1E 0657-558: Direct evidence for the existence of dark matter,” *Astrophys. J.*, vol. 604, no. 2, pp. 596—603, Apr 2004. [Online]. Available: <https://dx.doi.org/10.1086/381970>
- [19] M. Markevitch, A. H. Gonzalez, D. Clowe, A. Vikhlinin, W. Forman, C. Jones, S. Murray, and W. Tucker, “Direct constraints on the dark matter self-interaction cross section from the merging galaxy cluster 1E 0657-56,” *Astrophys. J.*, vol. 606, no. 2, pp. 819—824, May 2004. [Online]. Available: <https://dx.doi.org/10.1086/383178>
- [20] G. W. Angus, B. Famaey, and H. S. Zhao, “Can MOND take a bullet? Analytical comparisons of three versions of MOND beyond spherical symmetry,” *Mon. Not. R. Astron. Soc.*, vol. 371, no. 1, pp. 138—146, Aug 2006. [Online]. Available: <https://dx.doi.org/10.1111/j.1365-2966.2006.10668.x>
- [21] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, “A direct empirical proof of the existence of dark matter,” *Astrophys. J.*, vol. 648, no. 2, pp. L109—L113, Aug 2006. [Online]. Available: <https://dx.doi.org/10.1086/508162>
- [22] R. Massey, “Viewing dark matter with weak gravitational lensing from HST,” in *Problems of Practical Cosmology, Volume 1*, Y. V. Baryshev, I. N. Taganov, and P. Teerikorpi, Eds., vol. 1, Jan. 2008, pp. 41–49.
- [23] D. S. Graff and K. Freese, “Analysis of a Hubble Space Telescope search for red dwarfs: Limits on baryonic matter in the galactic halo,” *Astrophys. J.*, vol. 456, no. 1, Jan 1996. [Online]. Available: <https://dx.doi.org/10.1086/309850>
- [24] Tisserand, P., Le Guillou, L., Afonso, C., Albert, J. N., Andersen, J., Ansari, R., Aubourg, É., Baryshev, P., Beaulieu, J. P., Charlot, X. et al., “Limits on the macho content of the galactic halo from the EROS-2 survey of the Magellanic Clouds,” *Astron. Astrophys.*, vol. 469, no. 2, p. 387–404, Apr 2007. [Online]. Available: <https://dx.doi.org/10.1051/0004-6361:20066017>
- [25] G. Arcadi, M. Dutra, P. Ghosh, M. Lindner, Y. Mambrini, M. Pierre, S. Profumo, and F. S. Queiroz, “The waning of the WIMP? A review of models, searches, and

- constraints,” *Eur. Phys. J. C*, vol. 78, no. 3, p. 203, Mar 2018. [Online]. Available: <https://doi.org/10.1140/epjc/s10052-018-5662-y>
- [26] G. Jungman, M. Kamionkowski, and K. Griest, “Supersymmetric dark matter,” *Phys. Rep.*, vol. 267, no. 5–6, p. 195–373, Mar 1996. [Online]. Available: [https://dx.doi.org/10.1016/0370-1573\(95\)00058-5](https://dx.doi.org/10.1016/0370-1573(95)00058-5)
- [27] G. Steigman, B. Dasgupta, and J. F. Beacom, “Precise relic WIMP abundance and its impact on searches for dark matter annihilation,” *Phys. Rev. D*, vol. 86, no. 2, Jul 2012. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevD.86.023506>
- [28] D. A. Palken, “Modeling phonons in physical detectors,” Undergraduate Thesis, Bowdoin College, 2014.
- [29] M. Tanabashi, K. Hagiwara, K. Hikasa, K. Nakamura, Y. Sumino, F. Takahashi, J. Tanaka, K. Agashe, G. Aielli, C. AMSLER et al., “Review of particle physics,” *Phys. Rev. D*, vol. 98, p. 030001, Aug 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.98.030001>
- [30] D. S. Akerib, S. Alsum, H. M. Araújo, X. Bai, A. J. Bailey, J. Balajthy, P. Beltrame, E. P. Bernard, A. Bernstein, T. P. Biesiadzinski et al., “Results from a search for dark matter in the complete LUX exposure,” *Phys. Rev. Lett.*, vol. 118, p. 021303, Jan 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.118.021303>
- [31] R. Agnese, T. Aramaki, I. J. Arnuist, W. Baker, D. Balakishiyeva, S. Banik, D. Barker, R. Basu Thakur, D. A. Bauer, T. Binder et al., “Results from the Super Cryogenic Dark Matter Search experiment at Soudan,” *Phys. Rev. Lett.*, vol. 120, no. 6, Feb 2018. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevLett.120.061802>
- [32] E. Aprile, J. Aalbers, F. Agostini, M. Alfonsi, L. Althueser, F. D. Amaro, V. C. Antochi, E. Angelino, F. Arneodo, D. Barge et al., “Search for light dark matter interactions enhanced by the Migdal effect or Bremsstrahlung in XENON1T,” *Phys. Rev. Lett.*, vol. 123, no. 24, Dec 2019. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevLett.123.241803>
- [33] C. Boehm, D. Cerdeño, P. Machado, A. O.-D. Campo, and E. Reid, “How high is the neutrino floor?” *J. Cosmol. Astropart. Phys.*, vol. 2019, no. 01, p. 043, Jan 2019. [Online]. Available: <https://dx.doi.org/10.1088/1475-7516/2019/01/043>
- [34] C. A. O’Hare, A. M. Green, J. Billard, E. Figueroa-Feliciano, and L. E. Strigari, “Readout strategies for directional dark matter detection beyond the neutrino background,” *Phys. Rev. D*, vol. 92, no. 6, Sep 2015. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevD.92.063518>
- [35] L. Roszkowski, E. M. Sessolo, and S. Trojanowski, “WIMP dark matter candidates and searches—current status and future prospects,” *Rep. Prog. Phys.*, vol. 81, no. 6, p. 066201, May 2018. [Online]. Available: <https://dx.doi.org/10.1088/1361-6633/aab913>
- [36] R. D. Peccei and H. R. Quinn, “CP conservation in the presence of pseudoparticles,” *Phys. Rev. Lett.*, vol. 38, p. 1440, Jun 1977. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.38.1440>

- [37] R. D. Peccei and H. R. Quinn, “Constraints imposed by CP conservation in the presence of pseudoparticles,” *Phys. Rev. D*, vol. 16, p. 1791, Sep 1977. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.16.1791>
- [38] J. Schwinger, “The theory of quantized fields. I,” *Phys. Rev.*, vol. 82, pp. 914–927, Jun 1951. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.82.914>
- [39] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, “Evidence for the 2π decay of the k_2^0 meson,” *Phys. Rev. Lett.*, vol. 13, pp. 138–140, Jul 1964. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.13.138>
- [40] B. Aubert, D. Boutigny, J.-M. Gaillard, A. Hicheur, Y. Karyotakis, J. P. Lees, P. Robbe, V. Tisserand, A. Palano, G. P. Chen *et al.*, “Observation of CP violation in the B^0 meson system,” *Phys. Rev. Lett.*, vol. 87, no. 9, Aug 2001. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevLett.87.091801>
- [41] K. Abe, R. Akutsu, A. Ali, C. Alt, C. Andreopoulos, L. Anthony, M. Antonova, S. Aoki, A. Ariga, Y. Asada *et al.*, “Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations,” *Phys. Rev.*, vol. 580, pp. 339–344, Apr 2020. [Online]. Available: <https://doi.org/10.1038/s41586-020-2177-0>
- [42] B. M. Brubaker, “First results from the HAYSTAC axion search,” Ph.D. Thesis, Yale University, 2017.
- [43] C. Abel, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Burghoff, E. Chanel *et al.*, “Measurement of the permanent electric dipole moment of the neutron,” *Phys. Rev. Lett.*, vol. 124, p. 081803, Feb 2020. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.124.081803>
- [44] A. E. Nelson, “Naturally weak CP violation,” *Phys. Lett. B*, vol. 136, pp. 387–391, 1984.
- [45] S. M. Barr, “Solving the strong CP problem without the Peccei-Quinn symmetry,” *Phys. Rev. Lett.*, vol. 53, pp. 329–332, Jul 1984. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.53.329>
- [46] S. M. Barr, “Natural class of non-Peccei-Quinn models,” *Phys. Rev. D*, vol. 30, pp. 1805–1811, Oct 1984. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.30.1805>
- [47] P. Sikivie, “The pool-table analogy with axion physics,” *Phys. Today*, vol. 49, no. 12, p. 22–27, Dec 1996. [Online]. Available: <https://dx.doi.org/10.1063/1.881573>
- [48] F. Wilczek, “Time’s (almost) reversible arrow,” *Quanta magazine*, Jan 2016. [Online]. Available: <https://www.quantamagazine.org/how-axions-may-explain-times-arrow-20160107/>
- [49] S. Weinberg, “A new light boson?” *Phys. Rev. Lett.*, vol. 40, pp. 223–226, Jan 1978. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.40.223>
- [50] F. Wilczek, “Problem of strong P and T invariance in the presence of instantons,” *Phys. Rev. Lett.*, vol. 40, pp. 279–282, Jan 1978. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.40.279>

- [51] M. S. Seigar, “Cold dark matter, hot dark matter, and their alternatives,” in Dark Matter in the Universe, ser. 2053-2571. Morgan & Claypool Publishers, 2015, pp. 3–1 to 3–9. [Online]. Available: <https://dx.doi.org/10.1088/978-1-6817-4118-5ch3>
- [52] J. Preskill, M. B. Wise, and F. Wilczek, “Cosmology of the invisible axion,” Phys. Lett. B, vol. 120, p. 127, Jan 1983. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0370269383906378>
- [53] L. Abbott and P. Sikivie, “A cosmological bound on the invisible axion,” Phys. Lett. B, vol. 120, no. 1, p. 133, 1983. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/037026938390638X>
- [54] M. Dine and W. Fischler, “The not-so-harmless axion,” Phys. Lett. B, vol. 120, no. 1, pp. 137–141, 1983. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0370269383906391>
- [55] J. E. Kim, “Weak-interaction singlet and strong CP invariance,” Phys. Rev. Lett., vol. 43, p. 103, Jul 1979. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.43.103>
- [56] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, “Can confinement ensure natural CP invariance of strong interactions?” Nucl. Phys. B, vol. 166, no. 3, p. 493, 1980. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0550321380902096>
- [57] M. Dine, W. Fischler, and M. Srednicki, “A simple solution to the strong CP problem with a harmless axion,” Phys. Lett. B, vol. 104, no. 3, p. 199, 1981. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0370269381905906>
- [58] A. R. Zhitnitsky, “On possible suppression of the axion hadron interactions,” Sov. J. Nucl. Phys., vol. 31, p. 260, 1980.
- [59] S. L. Cheng, C. Q. Geng, and W.-T. Ni, “Axion-photon couplings in invisible axion models,” Phys. Rev. D, vol. 52, pp. 3132–3135, Sep 1995. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.52.3132>
- [60] I. G. Irastorza and J. Redondo, “New experimental approaches in the search for axion-like particles,” Prog. Part. Nucl. Phys., vol. 102, pp. 89–159, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S014664101830036X>
- [61] V. Anastassopoulos, S. Aune, K. Barth, A. Belov, H. Brauning, G. Cantatore, J. M. Carmona, J. F. Castel, S. A. Cetin, F. Christensen et al., “New CAST limit on the axion–photon interaction,” Nat. Phys., vol. 13, no. 6, pp. 584–590, May 2017. [Online]. Available: <https://dx.doi.org/10.1038/nphys4109>
- [62] M. Simanovskaia, “Design, fabrication, and characterization of a high-frequency microwave cavity for HAYSTAC,” Ph.D. Thesis, University of California Berkeley, 2019.
- [63] G. G. Raffelt, “Astrophysical axion bounds,” in Axions: Theory, Cosmology, and Experimental Searches. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 51–71. [Online]. Available: https://link.springer.com/10.1007/978-3-540-73518-2_{_}3
- [64] N. Bar, K. Blum, and G. D’Amico, “Is there a supernova bound on axions?” Phys. Rev. D, vol. 101, no. 12, Jun 2020. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevD.101.123025>

- [65] E. Berkowitz, M. I. Buchoff, and E. Rinaldi, “Lattice QCD input for axion cosmology,” *Phys. Rev. D*, vol. 92, p. 034507, Aug 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.92.034507>
- [66] S. Borsanyi, Z. Fodor, J. Guenther, K.-H. Kampert, S. Katz, T. Kawanai, T. Kovács, S. Mages, A. Pasztor, F. Pittler *et al.*, “Calculation of the axion mass based on high-temperature lattice quantum chromodynamics,” *Nature*, vol. 539, pp. 69–71, Nov 2016. [Online]. Available: <https://doi.org/10.1038/nature20115>
- [67] G. Ballesteros, J. Redondo, A. Ringwald, and C. Tamarit, “Unifying inflation with the axion, dark matter, baryogenesis, and the seesaw mechanism,” *Phys. Rev. Lett.*, vol. 118, p. 071802, Feb 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.118.071802>
- [68] G. Ballesteros, J. Redondo, A. Ringwald, and C. Tamarit, “Standard Model—axion—seesaw—Higgs portal inflation. Five problems of particle physics and cosmology solved in one stroke,” *J. Cosmol. Astropart. Phys.*, vol. 2017, no. 08, p. 001–001, Aug 2017. [Online]. Available: <https://dx.doi.org/10.1088/1475-7516/2017/08/001>
- [69] V. B. Klaer and G. D. Moore, “The dark-matter axion mass,” *J. Cosmol. Astropart. Phys.*, vol. 2017, no. 11, p. 049–049, Nov 2017. [Online]. Available: <https://dx.doi.org/10.1088/1475-7516/2017/11/049>
- [70] M. Buschmann, J. W. Foster, and B. R. Safdi, “Early-universe simulations of the cosmological axion,” *Phys. Rev. Lett.*, vol. 124, p. 161103, Apr 2020. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.124.161103>
- [71] J. I. Read, “The local dark matter density,” *J. Phys. G*, vol. 41, no. 6, p. 063101, May 2014. [Online]. Available: <https://dx.doi.org/10.1088/0954-3899/41/6/063101>
- [72] J. Herzog-Arbeitman, M. Lisanti, P. Madau, and L. Necib, “Empirical determination of dark matter velocities using metal-poor stars,” *Phys. Rev. Lett.*, vol. 120, no. 4, Jan 2018. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevLett.120.041102>
- [73] J. Herzog-Arbeitman, M. Lisanti, and L. Necib, “The metal-poor stellar halo in RAVE-TGAS and its implications for the velocity distribution of dark matter,” *J. Cosmol. Astropart. Phys.*, vol. 2018, no. 04, p. 052–052, Apr 2018. [Online]. Available: <https://dx.doi.org/10.1088/1475-7516/2018/04/052>
- [74] E. W. Lentz, T. R. Quinn, L. J. Rosenberg, and M. J. Tremmel, “A new signal model for axion cavity searches from N-body simulations,” *Astrophys. J.*, vol. 845, no. 2, p. 121, Aug 2017. [Online]. Available: <https://doi.org/10.3847/1538-4357/2017/845/2>
- [75] J. W. Foster, N. L. Rodd, and B. R. Safdi, “Revealing the dark matter halo with axion direct detection,” *Phys. Rev. D*, vol. 97, p. 123006, Jun 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.97.123006>
- [76] M. S. Turner, “Periodic signatures for the detection of cosmic axions,” *Phys. Rev. D*, vol. 42, p. 3572, Nov 1990. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.42.3572>
- [77] R. H. Dicke, “The measurement of thermal radiation at microwave frequencies,” *Rev. Sci. Instrum.*, vol. 17, no. 7, pp. 268–275, 1946. [Online]. Available: <https://doi.org/10.1063/1.1770483>

- [78] B. M. Brubaker, L. Zhong, S. K. Lamoreaux, K. W. Lehnert, and K. A. van Bibber, “HAYSTAC axion search analysis procedure,” *Phys. Rev. D*, vol. 96, p. 123008, Dec 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.96.123008>
- [79] S. Matsuki, I. Ogawa, S. Nakamura, M. Tada, K. Yamamoto, and A. Masaike, “Rydberg-atom cavity detector for dark matter axion search in Kyoto,” *Nucl. Phys. B*, vol. 51, no. 2, pp. 213–217, 1996, proceedings of the International Symposium on Sources and Detection of Dark Matter in the Universe. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S092056329600518X>
- [80] A. Kitagawa, K. Yamamoto, and S. Matsuki, “Quantum analysis of Rydberg atom cavity detector for dark matter axion search,” [arXiv preprint [arxiv:hep-ph/9908445](https://arxiv.org/abs/hep-ph/9908445), Aug 1999].
- [81] K. Yamamoto, M. Tada, Y. Kishimoto, M. Shibata, K. Kominato, T. Ooishi, S. Yamada, T. Saida, H. Funahashi, A. Masaike et al., “The Rydberg-atom-cavity axion search,” [arXiv preprint [arxiv:hep-ph/0101200](https://arxiv.org/abs/hep-ph/0101200), Jan 2001].
- [82] A. Dixit, A. Chou, and D. Schuster, “Detecting axion dark matter with superconducting qubits,” *Springer Proc. Phys.*, vol. 211, pp. 97–103, 2018.
- [83] S. K. Lamoreaux, K. A. van Bibber, K. W. Lehnert, and G. Carosi, “Analysis of single-photon and linear amplifier detectors for microwave cavity dark matter axion searches,” *Phys. Rev. D*, vol. 88, no. 3, Aug 2013. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevD.88.035020>
- [84] H. Zheng, M. Silveri, R. T. Brierley, S. M. Girvin, and K. W. Lehnert, “Accelerating dark-matter axion searches with quantum measurement technology,” [arXiv preprint [arxiv:1607.02529](https://arxiv.org/abs/1607.02529), Jul 2016].
- [85] P. Sikivie, “Experimental tests of the “invisible” axion,” *Phys. Rev. Lett.*, vol. 51, pp. 1415–1417, Oct 1983. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.51.1415>
- [86] P. Sikivie, “Detection rates for “invisible”-axion searches,” *Phys. Rev. D*, vol. 32, p. 2988, Dec 1985. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.32.2988>
- [87] H. Primakoff, “Photo-production of neutral mesons in nuclear electric fields and the mean life of the neutral meson,” *Phys. Rev.*, vol. 81, pp. 899–899, Mar 1951. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.81.899>
- [88] J. D. Jackson, *Classical electrodynamics*, 3rd ed. New York, NY: Wiley, 1999, pp. 352–406. [Online]. Available: <https://cdsweb.cern.ch/record/490457>
- [89] C. Hagmann, D. Kinion, W. Stoeffl, K. van Bibber, E. Daw, H. Peng, L. J. Rosenberg, J. LaVeigne, P. Sikivie, N. S. Sullivan et al., “Results from a high-sensitivity search for cosmic axions,” *Phys. Rev. Lett.*, vol. 80, p. 2043, Mar 1998. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.80.2043>
- [90] S. A. Kenany, M. A. Anil, K. M. Backes, B. M. Brubaker, S. B. Cahn, G. Carosi, Y. V. Gurevich, W. F. Kindel, S. K. Lamoreaux, K. W. Lehnert et al., “Design and operational experience of a microwave cavity axion detector for the 20 – 100 μeV range,” *Nucl. Instrum. Meth. A*, vol. 854, p. 11, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900217301948>

- [91] T. Braine, R. Cervantes, N. Crisosto, N. Du, S. Kimes, L. J. Rosenberg, G. Rybka, J. Yang, D. Bowring, A. S. Chou *et al.*, “Extended search for the invisible axion with the Axion Dark Matter Experiment,” *Phys. Rev. Lett.*, vol. 124, p. 101303, Mar 2020. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.124.101303>
- [92] M. Simanovskaia, A. Droster, H. Jackson, I. Urdinaran, and K. van Bibber, “A symmetric multi-rod tunable microwave cavity for the HAYSTAC dark matter axion search,” [arXiv preprint [arxiv:2006.01248](https://arxiv.org/abs/2006.01248), Jun 2020].
- [93] C. Boutan, M. Jones, B. H. LaRoque, N. S. Oblath, R. Cervantes, N. Du, N. Force, S. Kimes, R. Ottens, L. J. Rosenberg *et al.*, “Piezoelectrically tuned multimode cavity search for axion dark matter,” *Phys. Rev. Lett.*, vol. 121, p. 261302, Dec 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.121.261302>
- [94] B. T. McAllister, G. Flower, E. N. Ivanov, M. Goryachev, J. Bourhill, and M. E. Tobar, “The ORGAN experiment: An axion haloscope above 15 GHz,” *Phys. Dark Univ.*, vol. 18, p. 67, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2212686417300602>
- [95] S. Chaudhuri, K. Irwin, P. W. Graham, and J. Mardon, “Fundamental limits of electromagnetic axion and hidden-photon dark matter searches: Part I - the quantum limit,” [arXiv preprint [arxiv:1803.01627](https://arxiv.org/abs/1803.01627), Aug 2018].
- [96] N. M. Rapidis, S. M. Lewis, and K. A. van Bibber, “Characterization of the HAYSTAC axion dark matter search cavity using microwave measurement and simulation techniques,” *Rev. Sci. Instrum.*, vol. 90, no. 2, p. 024706, Feb 2019. [Online]. Available: <https://dx.doi.org/10.1063/1.5055246>
- [97] S. DePanfilis, A. C. Melissinos, B. E. Moskowitz, J. T. Rogers, Y. K. Semertzidis, W. U. Wuensch, H. J. Halama, A. G. Prodel, W. B. Fowler, and F. A. Nezrick, “Limits on the abundance and coupling of cosmic axions at $4.5 < m_a < 5.0 \mu\text{eV}$,” *Phys. Rev. Lett.*, vol. 59, p. 839, Aug 1987. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.59.839>
- [98] C. Hagmann, P. Sikivie, N. S. Sullivan, and D. B. Tanner, “Results from a search for cosmic axions,” *Phys. Rev. D*, vol. 42, p. 1297, Aug 1990. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.42.1297>
- [99] J. V. Sloan, M. Hotz, C. Boutan, R. Bradley, G. Carosi, D. Carter, J. Clarke, N. Crisosto, E. J. Daw, J. Gleason *et al.*, “Limits on axion-photon coupling or on local axion density: Dependence on models of the Milky Way’s dark halo,” *Phys. Dark Univ.*, vol. 14, p. 95, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2212686416300504>
- [100] W. U. Wuensch, S. DePanfilis-Wuensch, Y. K. Semertzidis, J. T. Rogers, A. C. Melissinos, H. J. Halama, B. E. Moskowitz, A. G. Prodel, W. B. Fowler, and F. A. Nezrick, “Results of a laboratory search for cosmic axions and other weakly coupled light particles,” *Phys. Rev. D*, vol. 40, p. 3153, Nov 1989. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.40.3153>
- [101] C. Hagmann, “A search for cosmic axions,” Ph.D. Thesis, University of Florida, 1990.

- [102] S. J. Asztalos, E. Daw, H. Peng, L. J. Rosenberg, D. B. Yu, C. Hagmann, D. Kinion, W. Stoeffl, K. van Bibber, J. LaVeigne *et al.*, “Experimental constraints on the axion dark matter halo density,” *Astrophys. J.*, vol. 571, no. 1, p. L27, 2002. [Online]. Available: <https://stacks.iop.org/1538-4357/571/i=1/a=L27>
- [103] S. J. Asztalos, R. F. Bradley, L. Duffy, C. Hagmann, D. Kinion, D. M. Moltz, L. J. Rosenberg, P. Sikivie, W. Stoeffl, N. S. Sullivan *et al.*, “Improved rf cavity search for halo axions,” *Phys. Rev. D*, vol. 69, p. 011101(R), Jan 2004. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.69.011101>
- [104] S. J. Asztalos, G. Carosi, C. Hagmann, D. Kinion, K. van Bibber, M. Hotz, L. J. Rosenberg, G. Rybka, J. Hoskins, J. Hwang *et al.*, “SQUID-based microwave cavity search for dark-matter axions,” *Phys. Rev. Lett.*, vol. 104, p. 041301, Jan 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.104.041301>
- [105] N. Du, N. Force, R. Khatiwada, E. Lentz, R. Ottens, L. J. Rosenberg, G. Rybka, G. Carosi, N. Woollett, D. Bowring *et al.*, “Search for invisible axion dark matter with the Axion Dark Matter Experiment,” *Phys. Rev. Lett.*, vol. 120, p. 151301, Apr 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.120.151301>
- [106] B. M. Brubaker, L. Zhong, Y. V. Gurevich, S. B. Cahn, S. K. Lamoreaux, M. Simanovskaia, J. R. Root, S. M. Lewis, S. Al Kenany, K. M. Backes *et al.*, “First results from a microwave cavity axion search at $24 \mu\text{eV}$,” *Phys. Rev. Lett.*, vol. 118, p. 061302, Feb 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.118.061302>
- [107] L. Zhong, S. Al Kenany, K. M. Backes, B. M. Brubaker, S. B. Cahn, G. Carosi, Y. V. Gurevich, W. F. Kindel, S. K. Lamoreaux, K. W. Lehnert *et al.*, “Results from phase 1 of the HAYSTAC microwave cavity axion experiment,” *Phys. Rev. D*, vol. 97, p. 092001, May 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.97.092001>
- [108] S. Lee, S. Ahn, J. Choi, B. R. Ko, and Y. K. Semertzidis, “Axion dark matter search around $6.7 \mu\text{eV}$,” *Phys. Rev. Lett.*, vol. 124, p. 101802, Mar 2020. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.124.101802>
- [109] D. Alesini, C. Braggio, G. Carugno, N. Crescini, D. D’Agostino, D. Di Gioacchino, R. Di Vora, P. Falferi, S. Gallo, U. Gambardella *et al.*, “Galactic axions search with a superconducting resonant cavity,” *Phys. Rev. D*, vol. 99, p. 101101(R), May 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.99.101101>
- [110] Y. Kahn, B. R. Safdi, and J. Thaler, “Broadband and resonant approaches to axion dark matter detection,” *Phys. Rev. Lett.*, vol. 117, p. 141801, Sep 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.117.141801>
- [111] J. L. Ouellet, C. P. Salemi, J. W. Foster, R. Henning, Z. Bogorad, J. M. Conrad, J. A. Formaggio, Y. Kahn, J. Minervini, A. Radovinsky *et al.*, “First results from ABRACADABRA-10 cm: A search for Sub- μeV axion dark matter,” *Phys. Rev. Lett.*, vol. 122, p. 121802, Mar 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.122.121802>
- [112] J. L. Ouellet, C. P. Salemi, J. W. Foster, R. Henning, Z. Bogorad, J. M. Conrad, J. A. Formaggio, Y. Kahn, J. Minervini, A. Radovinsky *et al.*, “Design and implementation of the

- ABRACADABRA-10 cm axion dark matter search,” *Phys. Rev. D*, vol. 99, no. 5, Mar 2019. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevD.99.052012>
- [113] G. Rybka, A. Wagner, K. Patel, R. Percival, K. Ramos, and A. Brill, “Search for dark matter axions with the Orpheus experiment,” *Phys. Rev. D*, vol. 91, p. 011701(R), Jan 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.91.011701>
- [114] P. Sikivie, N. Sullivan, and D. B. Tanner, “Proposal for axion dark matter detection using an *lc* circuit,” *Phys. Rev. Lett.*, vol. 112, p. 131301, Mar 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.112.131301>
- [115] N. Crisosto, P. Sikivie, N. S. Sullivan, D. B. Tanner, J. Yang, and G. Rybka, “ADMX SLIC: Results from a superconducting *LC* circuit investigating cold axions,” *Phys. Rev. Lett.*, vol. 124, p. 241101, Jun 2020. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.124.241101>
- [116] A. Caldwell, G. Dvali, B. Majorovits, A. Millar, G. Raffelt, J. Redondo, O. Reimann, F. Simon, and F. Steffen, “Dielectric haloscopes: A new way to detect axion dark matter,” *Phys. Rev. Lett.*, vol. 118, p. 091801, 2017. [Online]. Available: <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.118.091801>
- [117] A. J. Millar, G. G. Raffelt, J. Redondo, and F. D. Steffen, “Dielectric haloscopes to search for axion dark matter: Theoretical foundations,” *J. Cosmol. Astropart. Phys.*, vol. 2017, no. 01, p. 061–061, Jan 2017. [Online]. Available: <https://dx.doi.org/10.1088/1475-7516/2017/01/061>
- [118] K. Van Bibber, N. R. Dagdeviren, S. E. Koonin, A. K. Kerman, and H. N. Nelson, “Proposed experiment to produce and detect light pseudoscalars,” *Phys. Rev. Lett.*, vol. 59, pp. 759–762, Aug 1987. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.59.759>
- [119] K. Ehret, M. Frede, S. Ghazaryan, M. Hildebrandt, E.-A. Knabbe, D. Kracht, A. Lindner, J. List, T. Meier, N. Meyer *et al.*, “New ALPS results on hidden-sector lightweights,” *Phys. Lett. B*, vol. 689, no. 4, pp. 149–155, 2010. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0370269310005526>
- [120] R. Ballou, G. Deferne, M. Finger, M. Finger, L. Flekova, J. Hosek, S. Kunc, K. Macuchova, K. A. Meissner, P. Pugnati *et al.*, “New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall,” *Phys. Rev. D*, vol. 92, p. 092002, Nov 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.92.092002>
- [121] D. M. Lazarus, G. C. Smith, R. Cameron, A. C. Melissinos, G. Ruoso, Y. K. Semertzidis, and F. A. Nezrick, “Search for solar axions,” *Phys. Rev. Lett.*, vol. 69, pp. 2333–2336, Oct 1992. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.69.2333>
- [122] E. Armengaud, F. T. Avignone, M. Betz, P. Brax, P. Brun, G. Cantatore, J. M. Carmona, G. P. Carosi, F. Caspers, S. Caspi *et al.*, “Conceptual design of the International Axion Observatory (IAXO),” *J. Instrum.*, vol. 9, p. T05002, 2014.
- [123] A. Arvanitaki and A. A. Geraci, “Resonantly detecting axion-mediated forces with nuclear magnetic resonance,” *Phys. Rev. Lett.*, vol. 113, p. 161801, Oct 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.113.161801>

- [124] A. A. Geraci, H. Fosbinder-Elkins, C. Lohmeyer, J. Dargert, M. Cunningham, M. Harkness, E. Levenson-Falk, S. Mumford, A. Kapitulin, A. Arvanitaki *et al.*, “Progress on the ARIADNE axion experiment,” in *Microwave Cavities and Detectors for Axion Research*, G. Carosi, G. Rybka, and K. van Bibber, Eds., vol. 211. Springer International Publishing, 2018, pp. 151–161.
- [125] S. Asztalos, E. Daw, H. Peng, L. J. Rosenberg, C. Hagmann, D. Kinion, W. Stoeffl, K. van Bibber, P. Sikivie, N. S. Sullivan *et al.*, “Large-scale microwave cavity search for dark-matter axions,” *Phys. Rev. D*, vol. 64, p. 092003, Oct 2001. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.64.092003>
- [126] D. A. Palken, B. M. Brubaker, M. Malnou, S. A. Kenany, K. M. Backes, S. B. Cahn, Y. V. Gurevich, S. K. Lamoreaux, S. M. Lewis, R. H. Maruyama *et al.*, “Improved analysis framework for axion dark matter searches,” *Phys. Rev. D*, vol. 101, p. 123011, Jun 2020. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.101.123011>
- [127] A. B. Pippard, “The surface impedance of superconductors and normal metals at high frequencies II. The anomalous skin effect in normal metals,” *Proc. R. Soc. Lond. A*, vol. 191, pp. 385–399, Nov 1947. [Online]. Available: <https://doi.org/10.1098/rspa.1947.0122>
- [128] C. M. Caves, “Quantum limits on noise in linear amplifiers,” *Phys. Rev. D*, vol. 26, pp. 1817–1839, Oct 1982. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.26.1817>
- [129] H. W. Bode, *Network analysis and feedback amplifier design*. London: van Nostrand, Sep 1947.
- [130] R. M. Fano, “Theoretical limitations on the broadband matching of arbitrary impedances,” *J. Franklin. I.*, vol. 249, no. 1, Jan 1950. [Online]. Available: [https://doi.org/10.1016/0016-0032\(50\)90006-8](https://doi.org/10.1016/0016-0032(50)90006-8)
- [131] S. Chaudhuri, K. D. Irwin, P. W. Graham, and J. Mardon, “Optimal electromagnetic searches for axion and hidden-photon dark matter,” [arXiv preprint [arxiv:1904.05806](https://arxiv.org/abs/1904.05806), Apr 2019].
- [132] R. Gupta, M. Anerella, J. Cozzolino, P. Joshi, S. Joshi, S. Plate, W. Sampson, H. Song, P. Wanderer, W. Chung *et al.*, “Status of the 25 T, 100 mm bore HTS solenoid for an axion dark matter search experiment,” *IEEE Trans. Appl. Supercond.*, vol. 29, no. 5, pp. 1–5, Aug 2019.
- [133] I. Stern, A. A. Chisholm, J. Hoskins, P. Sikivie, N. S. Sullivan, D. B. Tanner, G. Carosi, and K. van Bibber, “Cavity design for high-frequency axion dark matter detectors,” *Rev. Sci. Instrum.*, vol. 86, no. 12, p. 123305, 2015. [Online]. Available: <https://doi.org/10.1063/1.4938164>
- [134] B. D. Josephson, “Possible new effects in superconductive tunnelling,” *Phys. Lett.*, vol. 1, no. 7, pp. 251–253, Jul. 1962.
- [135] P. W. Anderson and J. M. Rowell, “Probable observation of the Josephson superconducting tunneling effect,” *Phys. Rev. Lett.*, vol. 10, pp. 230–232, Mar 1963. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.10.230>
- [136] K. K. Likharev, *Dynamics of Josephson junctions and circuits*. Gordon and Breach science publishers, 1986.

- [137] B. Yurke, P. G. Kaminsky, R. E. Miller, E. A. Whittaker, A. D. Smith, A. H. Silver, and R. W. Simon, “Observation of 4.2-K equilibrium-noise squeezing via a Josephson-parametric amplifier,” *Phys. Rev. Lett.*, vol. 60, pp. 764–767, Feb 1988. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.60.764>
- [138] I. Siddiqi, R. Vijay, F. Pierre, C. M. Wilson, M. Metcalfe, C. Rigetti, L. Frunzio, and M. H. Devoret, “RF-driven Josephson bifurcation amplifier for quantum measurement,” *Phys. Rev. Lett.*, vol. 93, p. 207002, Nov 2004. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.93.207002>
- [139] M. A. Castellanos-Beltran and K. W. Lehnert, “Widely tunable parametric amplifier based on a superconducting quantum interference device array resonator,” *Appl. Phys. Lett.*, vol. 91, no. 8, p. 083509, 2007.
- [140] M. A. Castellanos-Beltran, K. D. Irwin, G. C. Hilton, L. R. Vale, and K. W. Lehnert, “Amplification and squeezing of quantum noise with a tunable Josephson metamaterial,” *Nat. Phys.*, vol. 4, no. 12, pp. 929–931, 2008. [Online]. Available: <https://dx.doi.org/10.1038/nphys1090>
- [141] T. Yamamoto, K. Inomata, M. Watanabe, K. Matsuba, T. Miyazaki, W. D. Oliver, Y. Nakamura, and J. S. Tsai, “Flux-driven Josephson parametric amplifier,” *Appl. Phys. Lett.*, vol. 93, no. 4, p. 042510, 2008. [Online]. Available: <https://dx.doi.org/10.1063/1.2964182>
- [142] N. Bergeal, F. Schackert, M. Metcalfe, R. Vijay, V. E. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret, “Phase-preserving amplification near the quantum limit with a Josephson ring modulator,” *Nature*, vol. 465, no. 7294, p. 64–68, May 2010. [Online]. Available: <https://dx.doi.org/10.1038/nature09035>
- [143] M. Hatridge, R. Vijay, D. H. Slichter, J. Clarke, and I. Siddiqi, “Dispersive magnetometry with a quantum limited SQUID parametric amplifier,” *Phys. Rev. B*, vol. 83, p. 134501, Apr 2011. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.83.134501>
- [144] A. Kamal, J. Clarke, and M. H. Devoret, “Gain, directionality, and noise in microwave SQUID amplifiers: Input-output approach,” *Phys. Rev. B*, vol. 86, p. 144510, Oct 2012. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.86.144510>
- [145] J. Y. Mutus, T. C. White, E. Jeffrey, D. Sank, R. Barends, J. Bochmann, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth *et al.*, “Design and characterization of a lumped element single-ended superconducting microwave parametric amplifier with on-chip flux bias line,” *Appl. Phys. Lett.*, vol. 103, no. 12, p. 122602, Sep 2013. [Online]. Available: <https://dx.doi.org/10.1063/1.4821136>
- [146] B. Abdo, A. Kamal, and M. Devoret, “Nondegenerate three-wave mixing with the Josephson ring modulator,” *Phys. Rev. B*, vol. 87, p. 014508, Jan 2013. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.87.014508>
- [147] J. Y. Mutus, T. C. White, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J. Kelly, A. Megrant *et al.*, “Strong environmental coupling in a Josephson parametric amplifier,” *Appl. Phys. Lett.*, vol. 104, no. 26, p. 263513, Jun 2014. [Online]. Available: <https://dx.doi.org/10.1063/1.4886408>

- [148] C. Eichler and A. Wallraff, “Controlling the dynamic range of a Josephson parametric amplifier,” *EPJ Quantum Technol.*, vol. 1, no. 1, p. 2, 2014. [Online]. Available: <https://doi.org/10.1140/epjqt2>
- [149] C. Macklin, K. O’Brien, D. Hover, M. E. Schwartz, V. Bolkhovskiy, X. Zhang, W. D. Oliver, and I. Siddiqi, “A near-quantum-limited Josephson traveling-wave parametric amplifier,” *Science*, 2015. [Online]. Available: <https://science.sciencemag.org/content/early/2015/09/02/science.aaa8525>
- [150] T. White, J. Mutus, I.-C. Hoi, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey et al., “Traveling wave parametric amplifier with Josephson junctions using minimal resonator phase matching,” *Appl. Phys. Lett.*, vol. 106, 03 2015.
- [151] F. Lecocq, L. Ranzani, G. A. Peterson, K. Cicak, R. W. Simmonds, J. D. Teufel, and J. Aumentado, “Nonreciprocal microwave signal processing with a field-programmable Josephson amplifier,” *Phys. Rev. Appl.*, vol. 7, p. 024028, Feb 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.7.024028>
- [152] N. E. Frattini, V. V. Sivak, A. Lingenfelter, S. Shankar, and M. H. Devoret, “Optimizing the nonlinearity and dissipation of a SNAIL parametric amplifier for dynamic range,” *Phys. Rev. Appl.*, vol. 10, no. 5, Nov 2018. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevApplied.10.054020>
- [153] V. V. Sivak, S. Shankar, G. Liu, J. Aumentado, and M. H. Devoret, “Josephson array-mode parametric amplifier,” *Phys. Rev. Appl.*, vol. 13, p. 024014, Feb 2020. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.13.024014>
- [154] R. Vijay, D. H. Slichter, and I. Siddiqi, “Observation of quantum jumps in a superconducting artificial atom,” *Phys. Rev. Lett.*, vol. 106, p. 110502, Mar 2011. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.106.110502>
- [155] J. D. Teufel, T. Donner, D. Li, J. W. Harlow, M. S. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, K. W. Lehnert, and R. W. Simmonds, “Sideband cooling of micromechanical motion to the quantum ground state,” *Nature*, vol. 475, no. 7356, p. 359–363, Jul 2011. [Online]. Available: <https://dx.doi.org/10.1038/nature10261>
- [156] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi, “Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback,” *Nature*, vol. 490, no. 7418, p. 77–80, Oct 2012. [Online]. Available: <https://dx.doi.org/10.1038/nature11505>
- [157] J. E. Johnson, C. Macklin, D. H. Slichter, R. Vijay, E. B. Weingarten, J. Clarke, and I. Siddiqi, “Heralded state preparation in a superconducting qubit,” *Phys. Rev. Lett.*, vol. 109, p. 050506, Aug 2012. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.109.050506>
- [158] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo, “Initialization by measurement of a superconducting quantum bit circuit,” *Phys. Rev. Lett.*, vol. 109, p. 050507, Aug 2012. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.109.050507>

- [159] P. Campagne-Ibarcq, E. Flurin, N. Roch, D. Darson, P. Morfin, M. Mirrahimi, M. H. Devoret, F. Mallet, and B. Huard, “Persistent control of a superconducting qubit by stroboscopic measurement feedback,” *Phys. Rev. X*, vol. 3, no. 2, May 2013. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevX.3.021008>
- [160] K. W. Murch, S. J. Weber, C. Macklin, and I. Siddiqi, “Observing single quantum trajectories of a superconducting quantum bit,” *Nature*, vol. 502, no. 7470, p. 211–214, Oct 2013. [Online]. Available: <https://dx.doi.org/10.1038/nature12539>
- [161] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. M. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin et al., “Quantum back-action of an individual variable-strength measurement,” *Science*, vol. 339, no. 6116, pp. 178–181, 2013.
- [162] H. S. Ku, W. F. Kindel, F. Mallet, S. Glancy, K. D. Irwin, G. C. Hilton, L. R. Vale, and K. W. Lehnert, “Generating and verifying entangled itinerant microwave fields with efficient and independent measurements,” *Phys. Rev. A*, vol. 91, no. 4, Apr 2015. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevA.91.042305>
- [163] W. F. Kindel, M. D. Schroer, and K. W. Lehnert, “Generation and efficient measurement of single photons from fixed-frequency superconducting qubits,” *Phys. Rev. A*, vol. 93, p. 033817, Mar 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.93.033817>
- [164] A. P. Reed, K. H. Mayer, J. D. Teufel, L. D. Burkhardt, W. Pfaff, M. Reagor, L. Sletten, X. Ma, R. J. Schoelkopf, E. Knill et al., “Faithful conversion of propagating quantum information to mechanical motion,” *Nat. Phys.*, vol. 13, no. 12, p. 1163–1167, Sep 2017. [Online]. Available: <https://dx.doi.org/10.1038/nphys4251>
- [165] J. J. Viennot, X. Ma, and K. W. Lehnert, “Phonon-number-sensitive electromechanics,” *Phys. Rev. Lett.*, vol. 121, p. 183601, Oct 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.121.183601>
- [166] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf et al., “Quantum error correction of a qubit encoded in grid states of an oscillator,” [arXiv preprint [arxiv:1907.12487](https://arxiv.org/abs/1907.12487), Jul 2019].
- [167] A. Eddins, J. Kreikebaum, D. Toyli, E. Levenson-Falk, A. Dove, W. Livingston, B. Levitan, L. Govia, A. Clerk, and I. Siddiqi, “High-efficiency measurement of an artificial atom embedded in a parametric amplifier,” *Phys. Rev. X*, vol. 9, no. 1, Jan 2019. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevX.9.011004>
- [168] L. R. Sletten, B. A. Moores, J. J. Viennot, and K. W. Lehnert, “Resolving phonon Fock states in a multimode cavity with a double-slit qubit,” *Phys. Rev. X*, vol. 9, p. 021056, Jun 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.9.021056>
- [169] X. Ma, J. J. Viennot, S. Kotler, J. D. Teufel, and K. W. Lehnert, “Nonclassical energy squeezing of a macroscopic mechanical oscillator,” [arXiv preprint [arxiv:2005.04260](https://arxiv.org/abs/2005.04260), May 2020].
- [170] E. I. Rosenthal, C. M. F. Schneider, M. Malnou, Z. Zhao, F. Leditzky, B. J. Chapman, W. Wustmann, X. Ma, D. A. Palken, M. Zanner et al., “Efficient and low-backaction quantum measurement using a chip-scale detector,” manuscript in preparation.

- [171] V. Ambegaokar and A. Baratoff, “Tunneling between superconductors,” *Phys. Rev. Lett.*, vol. 10, pp. 486–489, Jun 1963. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.10.486>
- [172] T. Van Duzer and C. W. Turner, *Principles of Superconductive Devices and Circuits*, 2nd ed. USA: Prentice Hall PTR, 1998.
- [173] M. A. Castellanos-Beltran, “Development of a Josephson parametric amplifier for the preparation and detection of nonclassical states of microwave fields,” Ph.D. Thesis, University of Colorado Boulder, 2010.
- [174] D. J. Griffiths and D. F. Schroeter, *Introduction to Quantum Mechanics*, 3rd ed. Cambridge University Press, 2018.
- [175] J. J. Sakurai and J. Napolitano, *Modern quantum mechanics*, 2nd ed. San Francisco, CA: Addison-Wesley, 2011, pp. 82–84. [Online]. Available: <https://cds.cern.ch/record/1341875>
- [176] B. Yurke, L. R. Corruccini, P. G. Kaminsky, L. W. Rupp, A. D. Smith, A. H. Silver, R. W. Simon, and E. A. Whittaker, “Observation of parametric amplification and deamplification in a Josephson parametric amplifier,” *Phys. Rev. A*, vol. 39, pp. 2519–2533, Mar 1989. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.39.2519>
- [177] S. Radic and C. McKinstrie, “Two-pump fiber parametric amplifiers,” *Opt. Fiber Technol.*, vol. 9, no. 1, pp. 7–23, 2003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S106852000200528X>
- [178] A. Kamal, A. Marblestone, and M. H. Devoret, “Signal-to-pump back action and self-oscillation in double-pump Josephson parametric amplifier,” *Phys. Rev. B*, vol. 79, p. 184301, May 2009. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.79.184301>
- [179] S. Boutin, D. M. Toyli, A. V. Venkatramani, A. W. Eddins, I. Siddiqi, and A. Blais, “Effect of higher-order nonlinearities on amplification and squeezing in Josephson parametric amplifiers,” *Phys. Rev. Appl.*, vol. 8, p. 054030, Nov 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.8.054030>
- [180] B. Yurke and E. Buks, “Performance of cavity-parametric amplifiers, employing Kerr nonlinearities, in the presence of two-photon loss,” *J. Lightwave Technol.*, vol. 24, no. 12, pp. 5054–5066, Dec 2006.
- [181] V. E. Manucharyan, E. Boaknin, M. Metcalfe, R. Vijay, I. Siddiqi, and M. H. Devoret, “Microwave bifurcation of a Josephson junction: Embedding-circuit requirements,” *Phys. Rev. B*, vol. 76, p. 014524, Jul 2007. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.76.014524>
- [182] T. Ojanen and J. Salo, “Possible scheme for on-chip element for squeezed microwave generation,” *Phys. Rev. B*, vol. 75, p. 184508, May 2007. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.75.184508>
- [183] L. Zhong, E. P. Menzel, R. D. Candia, P. Eder, M. Ihmig, A. Baust, M. Haeberlein, E. Hoffmann, K. Inomata, T. Yamamoto et al., “Squeezing with a flux-driven Josephson parametric amplifier,” *New J. Phys.*, vol. 15, no. 12, p. 125013, 2013. [Online]. Available: <https://stacks.iop.org/1367-2630/15/i=12/a=125013>

- [184] Z. R. Lin, K. Inomata, W. D. Oliver, K. Koshino, Y. Nakamura, J. S. Tsai, and T. Yamamoto, “Single-shot readout of a superconducting flux qubit with a flux-driven Josephson parametric amplifier,” *Appl. Phys. Lett.*, vol. 103, no. 13, p. 132602, 2013. [Online]. Available: <https://dx.doi.org/10.1063/1.4821822>
- [185] X. Zhou, V. Schmitt, P. Bertet, D. Vion, W. Wustmann, V. Shumeiko, and D. Esteve, “High-gain weakly nonlinear flux-modulated Josephson parametric amplifier using a SQUID array,” *Phys. Rev. B*, vol. 89, p. 214517, Jun 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.89.214517>
- [186] S. Pogorzalek, K. G. Fedorov, L. Zhong, J. Goetz, F. Wulschner, M. Fischer, P. Eder, E. Xie, K. Inomata, T. Yamamoto et al., “Hysteretic flux response and nondegenerate gain of flux-driven Josephson parametric amplifiers,” *Phys. Rev. Appl.*, vol. 8, p. 024012, Aug 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.8.024012>
- [187] D. M. Pozar, *Microwave Engineering*, 3rd ed. Hoboken, NJ: Wiley, 2005. [Online]. Available: <https://cds.cern.ch/record/882338>
- [188] H. S. Ku, “Generating and verifying entangled-itinerant microwave fields,” Ph.D. Thesis, University of Colorado Boulder, 2014.
- [189] C. W. Gardiner and M. J. Collett, “Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation,” *Phys. Rev. A*, vol. 31, pp. 3761–3774, Jun 1985. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.31.3761>
- [190] D. Walls and G. Milburn, *Quantum Optics*. Springer Berlin Heidelberg, 2008, pp. 127–141. [Online]. Available: <https://books.google.com/books?id=LiWsc3Nlf0kC>
- [191] J. Kerr, “XL. A new relation between electricity and light: Dielectrified media birefringent,” *Lond. Edinb. Dubl. Phil. Mag.*, vol. 50, no. 332, pp. 337–348, 1875. [Online]. Available: <https://doi.org/10.1080/14786447508641302>
- [192] J. Kerr, “LIV. A new relation between electricity and light: Dielectrified media birefringent (second paper),” *Lond. Edinb. Dubl. Phil. Mag.*, vol. 50, no. 333, pp. 446–458, 1875. [Online]. Available: <https://doi.org/10.1080/14786447508641319>
- [193] M. Scully and M. Zubairy, *Quantum Optics*. Cambridge University Press, 1997. [Online]. Available: <https://books.google.com/books?id=9lkgAwAAQBAJ>
- [194] C. Gerry and P. Knight, *Introductory Quantum Optics*. Cambridge University Press, 2004.
- [195] J. Thompson and H. Stewart, *Nonlinear Dynamics and Chaos*, 2nd ed., ser. Nonlinear Dynamics and Chaos. Wiley, 2002, p. 66. [Online]. Available: <https://books.google.com.ni/books?id=80ChNIpUDVAC>
- [196] A. Nayfeh and D. Mook, *Nonlinear Oscillations*, ser. Wiley Classics Library. Wiley, 2008, pp. 162–195. [Online]. Available: <https://books.google.com/books?id=sj3ebg7jRaoC>
- [197] W. Wustmann and V. Shumeiko, “Parametric resonance in tunable superconducting cavities,” *Phys. Rev. B*, vol. 87, p. 184501, May 2013. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.87.184501>

- [198] C. Eichler, “Experimental characterization of quantum microwave radiation and its entanglement with a superconducting qubit,” Ph.D. Thesis, ETH Zurich, 2013.
- [199] H. A. Haus and J. A. Mullen, “Quantum noise in linear amplifiers,” *Phys. Rev.*, vol. 128, pp. 2407–2413, Dec 1962. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.128.2407>
- [200] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, “Introduction to quantum noise, measurement, and amplification,” *Rev. Mod. Phys.*, vol. 82, p. 1155, Apr 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.82.1155>
- [201] A. T. Baust, “Characterization of flux-driven Josephson parametric amplifiers,” Ph.D. dissertation, Technische Universität München, 2010.
- [202] J. E. Sauvageau, P. A. Booi, M. W. Cromar, S. P. Benz, C. J. Burroughs, and J. A. Koch, “Superconducting integrated circuit fabrication process utilizing low temperature ECR-based PECVD SiO₂ dielectric films,” *IEEE Trans. Appl. Supercond.*, vol. 5, no. 2, 1995. [Online]. Available: <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=403046>
- [203] J. Mates, G. Hilton, K. Irwin, L. Vale, and K. Lehnert, “Demonstration of a multiplexer of dissipationless superconducting quantum interference devices,” *Appl. Phys. Lett.*, vol. 92, p. 023514, 2008-01 2008.
- [204] M. Malnou, D. A. Palken, L. R. Vale, G. C. Hilton, and K. W. Lehnert, “Optimal operation of a Josephson parametric amplifier for vacuum squeezing,” *Phys. Rev. Appl.*, vol. 9, p. 044023, Apr 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.9.044023>
- [205] J. Dunn, “EM ports — a ground is out there somewhere,” Mar 2011. [Online]. Available: <https://www.awr.com/articles/em-ports-ground-out-there-somewhere>
- [206] J. B. Johnson, “Thermal agitation of electricity in conductors,” *Phys. Rev.*, vol. 32, pp. 97–109, Jul 1928. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.32.97>
- [207] H. Nyquist, “Thermal agitation of electric charge in conductors,” *Phys. Rev.*, vol. 32, pp. 110–113, Jul 1928. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.32.110>
- [208] M. Malnou, D. A. Palken, B. M. Brubaker, L. R. Vale, G. C. Hilton, and K. W. Lehnert, “Squeezed vacuum used to accelerate the search for a weak classical signal,” *Phys. Rev. X*, vol. 9, p. 021023, May 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.9.021023>
- [209] K. W. Murch, S. J. Weber, K. M. Beck, E. Ginossar, and I. Siddiqi, “Reduction of the radiative decay of atomic coherence in squeezed vacuum,” *Nature*, vol. 499, no. 7456, p. 62–65, Jul 2013. [Online]. Available: <https://dx.doi.org/10.1038/nature12264>
- [210] S. Barzanjeh, D. P. DiVincenzo, and B. M. Terhal, “Dispersive qubit measurement by interferometry with parametric amplifiers,” *Phys. Rev. B*, vol. 90, p. 134515, Oct 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.90.134515>
- [211] N. Didier, A. Kamal, W. D. Oliver, A. Blais, and A. A. Clerk, “Heisenberg-limited qubit read-out with two-mode squeezed light,” *Phys. Rev. Lett.*, vol. 115, p. 093604, Aug 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.115.093604>

- [212] N. Didier, J. Bourassa, and A. Blais, “Fast quantum nondemolition readout by parametric modulation of longitudinal qubit-oscillator interaction,” *Phys. Rev. Lett.*, vol. 115, p. 203601, Nov 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.115.203601>
- [213] J. B. Clark, F. Lecocq, R. W. Simmonds, J. Aumentado, and J. D. Teufel, “Observation of strong radiation pressure forces from squeezed light on a mechanical oscillator,” *Nat. Phys.*, vol. 12, pp. 683–687, Nov 2016. [Online]. Available: <https://dx.doi.org/10.1038/nphys3701>
- [214] A. Bienfait, P. Campagne-Ibarcq, A. H. Küllerich, X. Zhou, S. Probst, J. J. Pla, T. Schenkel, D. Vion, D. Esteve, J. J. L. Morton *et al.*, “Magnetic resonance with squeezed microwaves,” *Phys. Rev. X*, vol. 7, p. 041011, Oct 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.7.041011>
- [215] F. Mallet, M. A. Castellanos-Beltran, H. S. Ku, S. Glancy, E. Knill, K. D. Irwin, G. C. Hilton, L. R. Vale, and K. W. Lehnert, “Quantum state tomography of an itinerant squeezed microwave field,” *Phys. Rev. Lett.*, vol. 106, p. 220502, Jun 2011. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.106.220502>
- [216] E. P. Menzel, R. Di Candia, F. Deppe, P. Eder, L. Zhong, M. Ihmig, M. Haeberlein, A. Baust, E. Hoffmann, D. Ballester *et al.*, “Path entanglement of continuous-variable quantum microwaves,” *Phys. Rev. Lett.*, vol. 109, p. 250502, Dec 2012. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.109.250502>
- [217] G. Liu, T.-C. Chien, X. Cao, O. Lanes, E. Alpern, D. Pekker, and M. Hatridge, “Josephson parametric converter saturation and higher order effects,” *Appl. Phys. Lett.*, vol. 111, no. 20, p. 202603, 2017. [Online]. Available: <https://doi.org/10.1063/1.5003032>
- [218] K. M. Sliwa, M. Hatridge, A. Narla, S. Shankar, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret, “Reconfigurable Josephson circulator/directional amplifier,” *Phys. Rev. X*, vol. 5, p. 041020, Nov 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.5.041020>
- [219] J. Kerckhoff, K. Lalumière, B. J. Chapman, A. Blais, and K. W. Lehnert, “On-chip superconducting microwave circulator from synthetic rotation,” *Phys. Rev. Appl.*, vol. 4, p. 034002, Sep 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.4.034002>
- [220] F. Lecocq, L. Ranzani, G. A. Peterson, K. Cicak, R. W. Simmonds, J. D. Teufel, and J. Aumentado, “Nonreciprocal microwave signal processing with a field-programmable Josephson amplifier,” *Phys. Rev. Appl.*, vol. 7, p. 024028, Feb 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.7.024028>
- [221] A. C. Mahoney, J. I. Colless, S. J. Pauka, J. M. Hornibrook, J. D. Watson, G. C. Gardner, M. J. Manfra, A. C. Doherty, and D. J. Reilly, “On-chip microwave quantum Hall circulator,” *Phys. Rev. X*, vol. 7, p. 011007, Jan 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.7.011007>
- [222] B. J. Chapman, E. I. Rosenthal, J. Kerckhoff, B. A. Moores, L. R. Vale, J. A. B. Mates, G. C. Hilton, K. Lalumière, A. Blais, and K. W. Lehnert, “Widely tunable on-chip microwave circulator for superconducting quantum circuits,” *Phys. Rev. X*, vol. 7, p. 041043, Nov 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.7.041043>

- [223] U. Leonhardt and H. Paul, “Measuring the quantum state of light,” *Prog. Quant. Electron.*, vol. 19, no. 2, pp. 89–130, 1995.
- [224] K. G. Fedorov, L. Zhong, S. Pogorzalek, P. Eder, M. Fischer, J. Goetz, E. Xie, F. Wulschner, K. Inomata, T. Yamamoto *et al.*, “Displacement of propagating squeezed microwave states,” *Phys. Rev. Lett.*, vol. 117, p. 020502, Jul 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.117.020502>
- [225] C. Eichler, Y. Salathe, J. Mlynek, S. Schmidt, and A. Wallraff, “Quantum-limited amplification and entanglement in coupled nonlinear resonators,” *Phys. Rev. Lett.*, vol. 113, p. 110502, Sep 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.113.110502>
- [226] R. Bradley, J. Clarke, D. Kinion, L. J. Rosenberg, K. van Bibber, S. Matsuki, M. Mück, and P. Sikivie, “Microwave cavity searches for dark-matter axions,” *Rev. Mod. Phys.*, vol. 75, pp. 777–817, Jun 2003. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.75.777>
- [227] C. M. Caves, “Quantum-mechanical noise in an interferometer,” *Phys. Rev. D*, vol. 23, pp. 1693–1708, Apr 1981. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.23.1693>
- [228] T. L. S. Collaboration, “A gravitational wave observatory operating beyond the quantum shot-noise limit,” *Nat. Phys.*, vol. 7, no. 12, p. 962, 2011. [Online]. Available: <https://dx.doi.org/10.1038/nphys2083>
- [229] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.*, “Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light,” *Nat. Photonics*, vol. 7, no. 8, p. 613, 2013. [Online]. Available: <https://dx.doi.org/10.1038/nphoton.2013.177>
- [230] H. J. Kimble, Y. Levin, A. B. Matsko, K. S. Thorne, and S. P. Vyatchanin, “Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics,” *Phys. Rev. D*, vol. 65, p. 022002, Dec 2001. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.65.022002>
- [231] M. Tse, H. Yu, N. Kijbunchoo, A. Fernandez-Galiana, P. Dupej, L. Barsotti, C. D. Blair, D. D. Brown, S. E. Dwyer, A. Effler *et al.*, “Quantum-enhanced Advanced LIGO detectors in the era of gravitational-wave astronomy,” *Phys. Rev. Lett.*, vol. 123, no. 23, p. 231107, Dec 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.123.231107>
- [232] P. Kurpiers, T. Walter, P. Magnard, Y. Salathe, and A. Wallraff, “Characterizing the attenuation of coaxial and rectangular microwave-frequency waveguides at cryogenic temperatures,” *EPJ Quantum Technol.*, vol. 4, no. 1, p. 8, May 2017. [Online]. Available: <https://doi.org/10.1140/epjqt/s40507-017-0059-7>
- [233] A. Savitzky and M. J. E. Golay, “Smoothing and differentiation of data by simplified least squares procedures.” *Anal. Chem.*, vol. 36, no. 8, p. 1627, 1964. [Online]. Available: <https://doi.org/10.1021/ac60214a047>

- [234] A. T. Malagon, “Search for 140 microeV pseudoscalar and vector dark matter using microwave cavities,” Ph.D. Thesis, Yale University, 2014.
- [235] L. Lista, Statistical Methods for Data Analysis in Particle Physics, 2nd ed. Springer, 2017.
- [236] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari et al., “Observation of gravitational waves from a binary black hole merger,” Phys. Rev. Lett., vol. 116, p. 061102, Feb 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>
- [237] G. Fragione, E. Grishin, N. W. C. Leigh, H. B. Perets, and R. Perna, “Black hole and neutron star mergers in galactic nuclei,” Mon. Not. R. Astron. Soc., vol. 488, no. 1, pp. 47–63, 06 2019. [Online]. Available: <https://doi.org/10.1093/mnras/stz1651>
- [238] S. B. McGrayne, The Theory That Would Not Die: How Bayes’ Rule Cracked the Enigma Code, Hunted Down Russian Submarines, and Emerged Triumphant from Two Centuries of Controversy. Yale University Press, 2011, p. 49. [Online]. Available: <https://www.jstor.org/stable/j.ctt1np76s>
- [239] D. Jackson, “CRISPR twins: Ethical time horizons,” ASEE Prism, vol. 28, no. 7, p. 21, Mar 2019.
- [240] R. L. Wasserstein and N. A. Lazar, “The ASA’s statement on p -values: Context, process, and purpose,” Am. Stat., vol. 70, no. 2, pp. 129–133, 2016. [Online]. Available: <https://doi.org/10.1080/00031305.2016.1154108>
- [241] F. Beaujean, A. Caldwell, and O. Reimann, “Is the bump significant? An axion-search example,” Eur. Phys. J. C, vol. 78, p. 793, 2018. [Online]. Available: <https://doi.org/10.1140/epjc/s10052-018-6217-y>
- [242] P. W. Graham, J. Mardon, and S. Rajendran, “Vector dark matter from inflationary fluctuations,” Phys. Rev. D, vol. 93, no. 10, May 2016. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevD.93.103520>
- [243] S. Chaudhuri, P. W. Graham, K. Irwin, J. Mardon, S. Rajendran, and Y. Zhao, “Radio for hidden-photon dark matter detection,” Phys. Rev. D, vol. 92, no. 7, Oct 2015. [Online]. Available: <https://dx.doi.org/10.1103/PhysRevD.92.075012>
- [244] M. Silva-Feaver, S. Chaudhuri, H. Cho, C. Dawson, P. Graham, K. Irwin, S. Kuenstner, D. Li, J. Mardon, H. Moseley et al., “Design overview of DM Radio pathfinder experiment,” IEEE Trans. Appl. Supercond., vol. 27, no. 4, p. 1, Jun 2017.
- [245] G. Shafer, A. Shen, N. Vereshchagin, and V. Vovk, “Test martingales, Bayes factors and p -values,” Stat. Sci., vol. 26, pp. 84–101, 2011.
- [246] Y. Zhang, S. Glancy, and E. Knill, “Asymptotically optimal data analysis for rejecting local realism,” Phys. Rev. A, vol. 84, p. 062118, Dec 2011. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.84.062118>
- [247] Y. Zhang, S. Glancy, and E. Knill, “Efficient quantification of experimental evidence against local realism,” Phys. Rev. A, vol. 88, p. 052119, Nov 2013. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.88.052119>

- [248] P. Grünwald, R. de Heide, and W. Koolen, “Safe testing,” [arXiv preprint [arxiv:1906.07801](https://arxiv.org/abs/1906.07801), Jun 2019].
- [249] P. W. Graham and S. Rajendran, “New observables for direct detection of axion dark matter,” *Phys. Rev. D*, vol. 88, p. 035023, Aug 2013. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.88.035023>
- [250] P. R. Bevington and D. K. Robinson, *Data reduction and error analysis for the physical sciences*, 3rd ed. McGraw-Hill, 2003, pp. 63–65.
- [251] T. S. Roussy, D. A. Palken, W. B. Cairncross, B. M. Brubaker, D. N. Gresh, M. Grau, K. C. Cossel, K. B. Ng, Y. Shagam, Y. Zhou *et al.*, “Experimental constraint on axion-like particle coupling over seven orders of magnitude in mass,” [arXiv preprint [arxiv:2006.15787](https://arxiv.org/abs/2006.15787), Jun 2020].
- [252] K. Wurtz, B. M. Brubaker, D. A. Palken, Y. Jiang, and K. W. Lehnert, “A cavity entanglement and state-swapping protocol to accelerate the search for axion dark matter,” manuscript in preparation.
- [253] K. M. Backes, D. A. Palken, S. Al Kenany, B. M. Brubaker, S. B. Cahn, A. Droster, G. C. Hilton, S. Ghosh, H. Jackson, S. K. Lamoreaux *et al.*, “A quantum-enhanced search for dark matter axions,” manuscript in preparation.
- [254] S. L. Braunstein and P. van Loock, “Quantum information with continuous variables,” *Rev. Mod. Phys.*, vol. 77, pp. 513–577, Jun 2005. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.77.513>
- [255] H. B. Callen and T. A. Welton, “Irreversibility and generalized noise,” *Phys. Rev.*, vol. 83, pp. 34–40, Jul 1951. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.83.34>
- [256] P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo, and A. Ringwald, “WISPy cold dark matter,” *J. Cosmol. Astropart. Phys.*, vol. 2012, no. 06, p. 013–013, Jun 2012. [Online]. Available: <https://dx.doi.org/10.1088/1475-7516/2012/06/013>
- [257] D. J. Marsh, “Axion cosmology,” *Phys. Rep.*, vol. 643, pp. 1–79, Jul 2016. [Online]. Available: <https://dx.doi.org/10.1016/j.physrep.2016.06.005><https://linkinghub.elsevier.com/retrieve/pii/S0370157316301557>
- [258] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, “String axiverse,” *Phys. Rev. D*, vol. 81, no. 12, p. 123530, Jun 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.81.123530>
- [259] W. Hu, R. Barkana, and A. Gruzinov, “Fuzzy cold dark matter: The wave properties of ultralight particles,” *Phys. Rev. Lett.*, vol. 85, no. 6, pp. 1158–1161, Aug 2000. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.85.1158>
- [260] A. Derevianko, “Detecting dark-matter waves with a network of precision-measurement tools,” *Phys. Rev. A*, vol. 97, no. 4, 2018.
- [261] H.-Y. Schive, T. Chiueh, and T. Broadhurst, “Cosmic structure as the quantum interference of a coherent dark wave,” *Nat. Phys.*, vol. 10, no. 7, pp. 496–499, Jul 2014. [Online]. Available: <https://www.nature.com/articles/nphys2996>

- [262] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten, “Ultralight scalars as cosmological dark matter,” *Phys. Rev. D*, vol. 95, no. 4, p. 043541, Feb 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.95.043541>
- [263] E. M. Purcell and N. F. Ramsey, “On the possibility of electric dipole moments for elementary particles and nuclei,” *Phys. Rev.*, vol. 78, no. 6, pp. 807–807, Jun 1950. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.78.807>
- [264] J. H. Smith, E. M. Purcell, and N. F. Ramsey, “Experimental limit to the electric dipole moment of the neutron,” *Phys. Rev.*, vol. 108, no. 1, pp. 120–122, Oct 1957. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.108.120>
- [265] J. M. Pendlebury, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, M. Burghoff, P. Geltenbort, K. Green et al., “Revised experimental upper limit on the electric dipole moment of the neutron,” *Phys. Rev. D*, vol. 92, p. 092003, Nov 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.92.092003>
- [266] J. J. Hudson, D. M. Kara, I. J. Smallman, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds, “Improved measurement of the shape of the electron,” *Nature*, vol. 473, no. 7348, pp. 493–496, May 2011. [Online]. Available: <https://dx.doi.org/10.1038/nature10104>
- [267] S. Eckel, P. Hamilton, E. Kirilov, H. W. Smith, and D. DeMille, “Search for the electron electric dipole moment using Ω -doublet levels in PbO,” *Phys. Rev. A*, vol. 87, p. 052130, May 2013. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.87.052130>
- [268] J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, I. Kozyryev et al., “Order of magnitude smaller limit on the electric dipole moment of the electron,” *Science*, vol. 343, no. 6168, pp. 269–272, Jan 2014. [Online]. Available: <https://www.sciencemag.org/cgi/doi/10.1126/science.1248213>
- [269] W. B. Cairncross, D. N. Gresh, M. Grau, K. C. Cossel, T. S. Roussy, Y. Ni, Y. Zhou, J. Ye, and E. A. Cornell, “Precision measurement of the electron’s electric dipole moment using trapped molecular ions,” *Phys. Rev. Lett.*, vol. 119, p. 153001, Oct 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.119.153001>
- [270] ACME Collaboration, “Improved limit on the electric dipole moment of the electron,” *Nature*, vol. 562, no. 7727, pp. 355–360, 2018. [Online]. Available: <https://www.nature.com/articles/s41586-018-0599-8>
- [271] D. Budker, P. W. Graham, M. Ledbetter, S. Rajendran, and A. O. Sushkov, “Proposal for a Cosmic Axion Spin Precession Experiment (CASPER),” *Phys. Rev. X*, vol. 4, p. 021030, May 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.4.021030>
- [272] P. W. Graham and S. Rajendran, “Axion dark matter detection with cold molecules,” *Phys. Rev. D*, vol. 84, no. 5, p. 055013, Sep 2011. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.84.055013>
- [273] C. Abel, N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Daum, M. Fairbairn, V. V. Flambaum, P. Geltenbort et al., “Search for axionlike dark matter through nuclear spin precession in electric and magnetic fields,” *Phys. Rev. X*, vol. 7, no. 4, pp. 1–9, 2017.

- [274] C. T. Hill, “Axion induced oscillating electric dipole moments,” *Phys. Rev. D*, vol. 91, no. 11, p. 111702, Jun 2015. [Online]. Available: <https://doi.org/10.1103/PhysRevD.91.111702><https://link.aps.org/doi/10.1103/PhysRevD.91.111702>
- [275] Y. V. Stadnik and V. V. Flambaum, “Axion-induced effects in atoms, molecules, and nuclei: Parity nonconservation, anapole moments, electric dipole moments, and spin-gravity and spin-axion momentum couplings,” *Phys. Rev. D*, vol. 89, p. 043522, Feb 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.89.043522>
- [276] K. C. Cossel, D. N. Gresh, L. C. Sinclair, T. Coffey, L. V. Skripnikov, A. N. Petrov, N. S. Mosyagin, A. V. Titov, R. W. Field, E. R. Meyer *et al.*, “Broadband velocity modulation spectroscopy of HfF^+ : Towards a measurement of the electron electric dipole moment,” *Chem. Phys. Lett.*, vol. 546, pp. 1–11, 2012. [Online]. Available: <https://dx.doi.org/10.1016/j.cplett.2012.06.037>
- [277] H. Loh, K. C. Cossel, M. C. Grau, K.-K. Ni, E. R. Meyer, J. L. Bohn, J. Ye, and E. A. Cornell, “Precision spectroscopy of polarized molecules in an ion trap,” *Science*, vol. 342, no. 6163, pp. 1220–1222, Dec 2013. [Online]. Available: <https://www.sciencemag.org/cgi/doi/10.1126/science.1243683>
- [278] M. Pospelov and I. Khriplovich, “Electric dipole moment of the W boson and the electron in the Kobayashi-Maskawa model,” *Sov. J. Nucl. Phys.*, vol. 53, pp. 638–640, 1991.
- [279] M. Pospelov and A. Ritz, “Electric dipole moments as probes of new physics,” *Ann. Phys.*, vol. 318, no. 1, pp. 119–169, Jul 2005. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0003491605000539>
- [280] M. S. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko, and C. W. Clark, “Search for new physics with atoms and molecules,” *Rev. Mod. Phys.*, vol. 90, p. 025008, Jun 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.90.025008>
- [281] Y. Nakai and M. Reece, “Electric dipole moments in natural supersymmetry,” *J. High Energy Phys.*, vol. 2017, no. 8, p. 31, Aug 2017. [Online]. Available: [https://link.springer.com/10.1007/JHEP08\(2017\)031](https://link.springer.com/10.1007/JHEP08(2017)031)
- [282] J. Engel, M. J. Ramsey-Musolf, and U. Van Kolck, “Electric dipole moments of nucleons, nuclei, and atoms: The Standard Model and beyond,” *Prog. Part. Nucl. Phys.*, vol. 71, pp. 21–74, 2013. [Online]. Available: <https://dx.doi.org/10.1016/j.pnpnp.2013.03.003>
- [283] N. F. Ramsey, “A molecular beam resonance method with separated oscillating fields,” *Phys. Rev.*, vol. 78, no. 6, pp. 695–699, Jun 1950.
- [284] W. Paul, “Electromagnetic traps for charged and neutral particles,” *Rev. Mod. Phys.*, vol. 62, pp. 531–540, Jul 1990. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.62.531>
- [285] E. R. Meyer, A. E. Leanhardt, E. A. Cornell, and J. L. Bohn, “Berry-like phases in structured atoms and molecules,” *Phys. Rev. A*, vol. 80, p. 062110, Dec 2009. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.80.062110>

- [286] V. V. Flambaum, B. M. Roberts, and Y. V. Stadnik, “Comment on “Axion induced oscillating electric dipole moments”,” *Phys. Rev. D*, vol. 95, p. 058701, Mar 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.95.058701>
- [287] Flambaum, V. V. (unpublished calculations).
- [288] V. V. Flambaum and H. B. T. Tan, “Oscillating nuclear electric dipole moment induced by axion dark matter produces atomic and molecular electric dipole moments and nuclear spin rotation,” *Phys. Rev. D*, vol. 100, p. 111301, Dec 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.100.111301>
- [289] V. V. Flambaum, M. Pospelov, A. Ritz, and Y. V. Stadnik, “Sensitivity of EDM experiments in paramagnetic atoms and molecules to hadronic CP violation,” [arXiv preprint [arxiv:1912.13129](https://arxiv.org/abs/1912.13129), Dec 2019].
- [290] J. Taylor and S. Hamilton, “Some tests of the Vaníček Method of spectral analysis,” *Astrophys. Space Sci.*, vol. 17, no. 2, pp. 357–367, Aug. 1972.
- [291] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. USA: Cambridge University Press, 2007, pp. 685–692.
- [292] G. P. Centers, J. W. Blanchard, J. Conrad, N. L. Figueroa, A. Garcon, A. V. Gramolin, D. F. J. Kimball, M. Lawson, B. Pelssers, J. A. Smiga *et al.*, “Stochastic amplitude fluctuations of bosonic dark matter and revised constraints on linear couplings,” [arXiv preprint [arxiv:1905.13650](https://arxiv.org/abs/1905.13650), May 2019].
- [293] A. Garcon, J. W. Blanchard, G. P. Centers, N. L. Figueroa, P. W. Graham, D. F. Jackson Kimball, S. Rajendran, A. O. Sushkov, Y. V. Stadnik, A. Wickenbrock *et al.*, “Constraints on bosonic dark matter from ultralow-field nuclear magnetic resonance,” *Sci. Adv.*, vol. 5, no. 10, 2019. [Online]. Available: <https://advances.sciencemag.org/content/5/10/eaax4539>
- [294] T. Wu, J. W. Blanchard, G. P. Centers, N. L. Figueroa, A. Garcon, P. W. Graham, D. F. J. Kimball, S. Rajendran, Y. V. Stadnik, A. O. Sushkov *et al.*, “Search for axionlike dark matter with a liquid-state nuclear spin comagnetometer,” *Phys. Rev. Lett.*, vol. 122, p. 191302, May 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.122.191302>
- [295] W. A. Terrano, E. G. Adelberger, C. A. Hagedorn, and B. R. Heckel, “Constraints on axionlike dark matter with masses down to 10^{-23} eV/ c^2 ,” *Phys. Rev. Lett.*, vol. 122, p. 231301, Jun 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.122.231301>
- [296] K. Van Tilburg, N. Leefer, L. Bougas, and D. Budker, “Search for ultralight scalar dark matter with atomic spectroscopy,” *Phys. Rev. Lett.*, vol. 115, p. 011802, Jun 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.115.011802>
- [297] A. Hees, J. Guéna, M. Abgrall, S. Bize, and P. Wolf, “Searching for an oscillating massive scalar field as a dark matter candidate using atomic hyperfine frequency comparisons,” *Phys. Rev. Lett.*, vol. 117, p. 061301, Aug 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.117.061301>

- [298] V. V. Flambaum, I. B. Samsonov, and H. B. Tran Tan, “Limits on CP -violating hadronic interactions and proton EDM from paramagnetic molecules,” [arXiv preprint [arxiv:2004.10359](https://arxiv.org/abs/2004.10359), Apr 2020].
- [299] K. Blum, R. T. D’Agnolo, M. Lisanti, and B. R. Safdi, “Constraining axion dark matter with Big Bang Nucleosynthesis,” *Phys. Lett. B*, vol. 737, pp. 30–33, 2014. [Online]. Available: <https://dx.doi.org/10.1016/j.physletb.2014.07.059>
- [300] G. G. Raffelt, “Astrophysical methods to constrain axions and other novel particle phenomena,” *Phys. Rep.*, vol. 198, no. 1–2, pp. 1–113, 1990.
- [301] R. E. Slusher, L. W. Hollberg, B. Yurke, J. C. Mertz, and J. F. Valley, “Observation of squeezed states generated by four-wave mixing in an optical cavity,” *Phys. Rev. Lett.*, vol. 55, pp. 2409–2412, Nov 1985. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.55.2409>
- [302] C. Boutan and ADMX Collaboration, “The Axion Dark Matter Experiment (ADMX) 4-cavity array,” in *APS April Meeting Abstracts*, ser. APS Meeting Abstracts, vol. 2019, Jan 2019, p. T04.006.
- [303] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, Aug 2018. [Online]. Available: <https://dx.doi.org/10.22331/q-2018-08-06-79>
- [304] F. Arute, K. Arya, R. Babbush, D. Bacon, J. Bardin, R. Barends, R. Biswas, S. Boixo, F. Brandao, D. Buell et al., “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, p. 505–510, 2019. [Online]. Available: <https://www.nature.com/articles/s41586-019-1666-5>
- [305] A. Narla, K. M. Sliwa, M. Hatridge, S. Shankar, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret, “Wireless Josephson amplifier,” *Appl. Phys. Lett.*, vol. 104, no. 23, p. 232605, Jun 2014. [Online]. Available: <https://dx.doi.org/10.1063/1.4883373>
- [306] B. Abdo, N. T. Bronn, O. Jinka, S. Olivadese, M. Brink, and J. M. Chow, “Multi-path interferometric Josephson directional amplifier for qubit readout,” *Quantum Sci. Technol.*, vol. 3, no. 2, p. 024003, Jan 2018. [Online]. Available: <https://iopscience.iop.org/article/10.1088/2058-9565/aaa1f9>
- [307] A. J. Sirois, M. A. Castellanos-Beltran, M. P. DeFeo, L. Ranzani, F. Lecocq, R. W. Simmonds, J. D. Teufel, and J. Aumentado, “Coherent-state storage and retrieval between superconducting cavities using parametric frequency conversion,” *Appl. Phys. Lett.*, vol. 106, no. 17, p. 172603, Apr 2015. [Online]. Available: <https://dx.doi.org/10.1063/1.4919759>
- [308] T.-C. Chien, O. Lanes, C. Liu, X. Cao, P. Lu, S. Motz, G. Liu, D. Pekker, and M. Hatridge, “Multiparametric amplification and qubit measurement with a Kerr-free Josephson ring modulator,” *Phys. Rev. A*, vol. 101, p. 042336, Apr 2020. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.101.042336>
- [309] A. Metelmann and A. A. Clerk, “Nonreciprocal photon transmission and amplification via reservoir engineering,” *Phys. Rev. X*, vol. 5, p. 021025, Jun 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.5.021025>

- [310] B. J. Chapman, “Widely tunable on-chip microwave circulator for superconducting quantum circuits,” Ph.D. Thesis, University of Colorado Boulder, 2017.

Appendix A

NIST trilayer fabrication process

The Josephson parametric amplifiers (JPAs) used for the work in this thesis were fabricated at NIST Boulder by Leila Vale using the NIST Nb-AlO_x-Nb trilayer process [202], modified by removing the shunt resistor layer and minimizing deposited oxides [203]. JPAs, which reside on individual chips, are produced in bulk on wafers which are then diced. The fabrication process is represented graphically in Fig. A.1.

Figure A.1a–d show the steps used to fabricate the Josephson junctions. First, Figure A.1a shows the three layers²⁰⁷ deposited uniformly across the entire wafer, which is made of Si with relative dielectric constant $\epsilon_r = 11.45$ [310] plus a thin oxide layer. The layers are deposited within a dedicated load-locked vacuum system. A Nb layer on the bottom forms the “base electrode” (BE), which most of the superconducting circuit wiring will be made of. Atop this, a film of Al is deposited and then oxidized. This AlO₂ oxide is the insulating layer at the center of the Josephson junction, Fig. 3.1. Finally, a second “counter electrode” (CE) layer of Nb sits atop the AlO₂. The circuitry is then formed subtractively, by etching away the CE, Al–AlO₂, and finally BE layers. What is left are just BE-connected pairs of Josephson junctions, Fig. A.1e, plus any BE wiring used to form transmission lines, interdigitated finger capacitors, etc.

Fig. A.1f–h show the remaining steps performed to connect the junctions. An electron cyclotron resonance plasma enhanced chemical vapor deposition process adds an insulating layer of

²⁰⁷ The three layers in the “trilayer” name refer to the layers that comprise the superconductor, insulator, and superconductor of the Josephson junction, Fig. 3.1. These are the two layers of Nb in Fig A.1a, and the aluminum oxide between them. The Al exists here in order to accommodate its oxide; it is also a superconductor at the temperatures at which we operate, so it can be thought of as part of the first superconductor.

SiO₂, which is selectively etched away to form vias above the junctions. Now, a wiring layer of Nb (called the wiring 1, or W1 layer) is added, forming the superconducting connections between pairs of junctions shown in Fig. A.1g, before etching excess SiO₂ away from the junctions in order to minimize dielectric loss.

The result, Fig. A.1h, is a pair of SQUIDs. Current entering from the left can flow easily through the superconductor. It splits into two parallel paths, tunneling up through the junctions, then left-to-right across the top layer of Nb, and finally down through the second pair of junctions. Because our arrays always start and end with the BE layer in order to connect to the rest of the circuitry, SQUIDs for us always come in pairs (junctions in fours): the cell shown in Fig. A.1h is simply reproduced $N/2$ times to make an array of N SQUIDs. For a complete set of design rules useful for laying out superconducting circuits to be fabricated with this process, see Ref. [310].

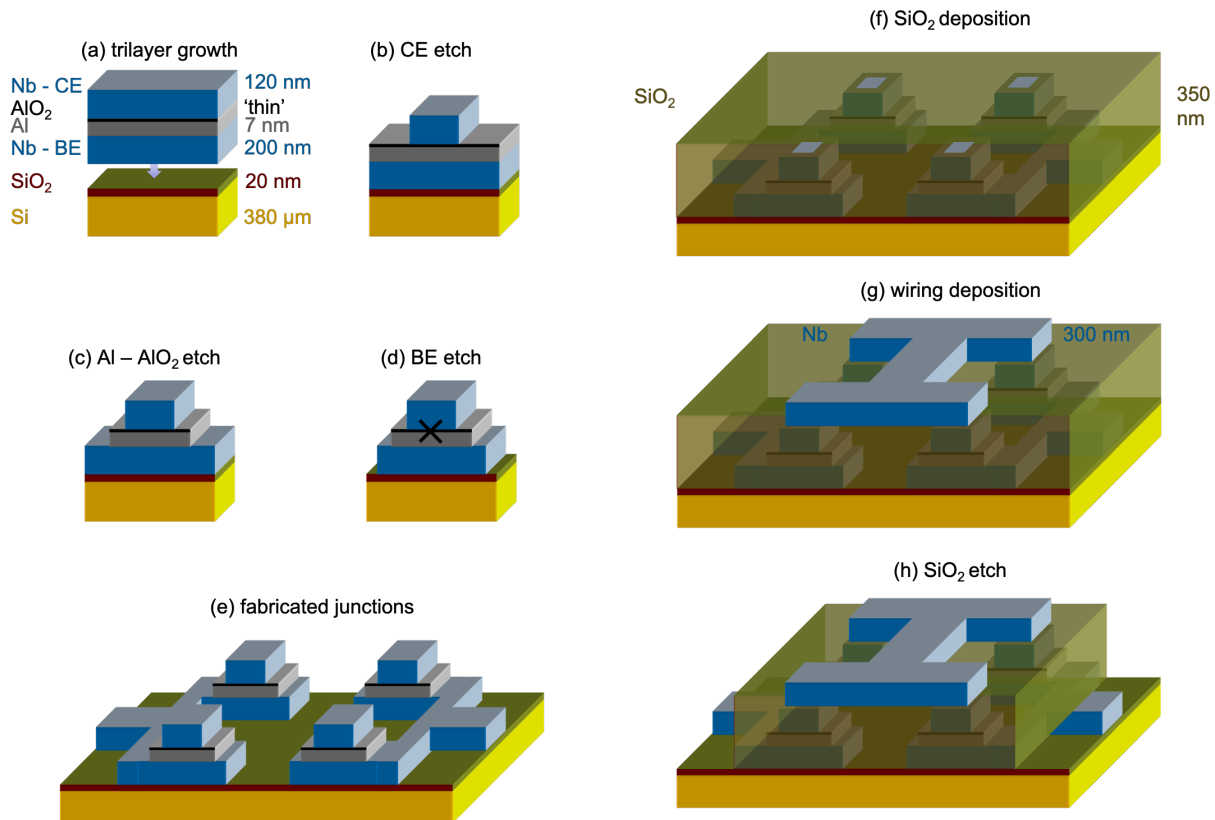


Figure A.1: (a) Trilayer growth. The layers of base electrode (BE) Nb, AlO₂ grown from Al, and counter electrode (CE) Nb, sit atop a Si wafer with a thin layer of oxide. Typical thicknesses are shown; for relevant lengths and widths, see Fig. 3.16. (b) CE etch. The CE is etched off away from the junctions. (c) Al-AIO₂ etch. The same process is used to remove the Al and its oxide. (d) BE etch. BE is removed anywhere it is not needed for junctions or other circuit wiring. The junction is marked with an “X” symbol. (e) Fabricated junctions. A cell of four junctions, which are not yet connected, remains. (f) SiO₂ deposition. A dielectric layer of SiO₂ is deposited. (g) wiring deposition. The SiO₂ is etched away above the junctions, forming vias. A wiring layer of Nb is then deposited on top, electrically connecting junctions to form pairs of SQUIDs. (h) SiO₂ etch. Undesired SiO₂ is removed with a final etch. Figure not to scale. Graphics courtesy of Maxime Malnou.

Appendix B

Resolution to the flux-tuning problem for flux-pumped JPAs

The flux-pumped Josephson parametric amplifiers (JPAs) used throughout this thesis worked generally well, but for the problem with their flux tuning that is evident in Fig. 3.32, especially as compared to the cleaner case of our current-pumped JPAs, Fig. 3.22. Every flux-pumped JPA we cooled down had some variant of this problem, though the specific structure varied from device to device, and sometimes upon thermally cycling the dilution refrigerator.




Figure B.1: Location of cuts made on dummy SQUIDs. The dummy SQUIDs are the blue-hashed region at the end of the SQUID array in this design layout. The pink hashed region above and below the SQUID array is the flux-bias line. In order to resolve the flux-tuning problem for our flux-pumped Josephson parametric amplifiers (JPAs), we have the Nb base electrode layer cut with a focused ion beam (FIB) at NIST. Any two of the three cuts that were made (red, dashed lines) on each pair of dummy SQUIDs is sufficient to break superconducting loops while maintaining electrical continuity to the SQUIDs.

The culprit turned out to be pairs of “dummy SQUIDs” that were placed at each end of the array. The purpose of the dummy SQUIDs, shown in Fig. B.1, was to mimic the superconducting metalization pattern of the actual SQUIDs. The aim was for SQUIDs at the end of the arrays not to see a different magnetic flux from those towards the center. Instead, the dummy SQUIDs, which

formed superconducting loops, readily trapped integer numbers of flux quanta when the devices were cooled down.

The dummy SQUIDs were revealed to be the issue when cutting them solved the problem. David Rudman at NIST Boulder performed the three cuts indicated in Fig. B.1 with a focused ion beam (FIB) for the dummy SQUIDs on both ends of the array, for two nominal-4.6 GHz devices. The cuts were performed so as to remove the superconducting loops but preserve electrical continuity on both sides of the array.




Figure B.2: (a), (b) Bidirectional flux-tuning sweep of FIB-cut JPA showing fixed tuning problem. Reflection coefficient phase $\arg(S_{11})$ is obtained via a flux-tuning sweep of DC bias current and probe frequency ν such as that in Fig. 3.32, with the coil current being swept up (a) and down (b). The flux-tuning curves no longer have the structure observed in Fig. 3.32. More than 1 GHz below the bare resonance, the bidirectional sweep reveals some hysteresis.

Two flux-tuning curves taken on one of the FIB-cut flux-pumped JPAs are shown in Fig. B.2. Figure B.2a shows the result of sweeping the coil current upward, and Fig. B.2b sweeping downward. The flux tuning is much improved from that of the unmodified flux-pumped JPAs, Fig. 3.32, strongly indicating that the superconducting loops formed by the dummy SQUIDs were causing the problem. In particular, there is far less structure, the zero-current resonance sits closer to zero flux, and the bare resonance frequency rose 0.5–1 GHz above the nominal value for the two chips measured. All

three effects can be explained by flux not being trapped near the SQUIDs.²⁰⁸

There is, however, a sharp feature that appears in Fig. B.2a at close to 3.5 mA of DC bias current. The same feature occurs at the mirrored location in Fig. B.2b. Operation near this hysteretic feature proved unreliable, but operation anywhere above it in resonance frequency was excellent.




Figure B.3: (a) Two-quadrature gain map in the two-dimensional parameter space of half-pump frequency $\nu_p/2$ and pump power estimated at JPA input. The white line shows the frequency of maximum gain at each power up until the region where oscillation occurs. (b) Gain G_s versus probe frequency ν traces taken at each point along the white line in (a), with colors matching those on the gain map.

Figure B.3 shows operation of the second FIB-cut flux-pumped JPA for gain at a bare resonance near 4.8 GHz. The gain map (Fig. B.3a) is clean, with traces shown up to 35 dB (Fig. B.3b), before oscillation sets in.

The results shown in Fig. B.2 for the first FIB-cut JPA and Fig. B.3 for the second were very similar across the two devices. The improved flux-tuning structure, zero-current behavior, and bare resonance, along with the robust amplifier operation, and the hysteresis, were visible in both. The only noticeable difference was in the specific value of the bare resonant frequency, 0.5 GHz higher for the second device. Despite the hysteresis, whose cause is as yet unknown, these

²⁰⁸ It is not surprising that the resonance actually moved above the nominal value absent trapped flux, as the critical currents for most of our flux-pumped JPAs were observed to be $\sim 20\%$ high.

devices represent an improvement over the flux-pumped JPAs used in this thesis. As such, future fabrication runs will not include dummy SQUIDs which form superconducting loops.

Appendix C

Effect of improper tuning rod thermalization on scan rate

The HAYSTAC Phase II measurement presented in Ch. 8 of this thesis is adversely affected by an improperly thermalized tuning rod. During Phase I operation of HAYSTAC, the issue was nicknamed the “hot rod” problem, for the tuning rod’s tendency to maintain a higher physical temperature than the base temperature of the fridge [42]. Between Runs I [106] and II [107] of Phase I, the hot rod problem was partially mitigated at some cost to cavity quality factor,²⁰⁹ but did not go away entirely. In future runs of HAYSTAC, the thermalization problem will be addressed further, but, in the context of the results presented in Sec. 8.3, it is important to understand the effect of the hot rod on (1) the absolute scan rate and (2) the improvement available from squeezing. We will find with respect to (1) that the hot rod reduces absolute scan rate by almost a factor of two. With respect to (2), we will see that it has virtually no effect on the scan rate enhancement available from squeezing. Indeed, since the reflected noise is still at the quantum limit, it is still the quantum noise that is being squeezed.

To compare the effects of hot rod versus cold rod (i.e. rod at fridge temperature, $S_c = S_f$) when squeezing versus not squeezing, we first compute the noise and signal power spectral densities for all four cases using the model represented in Fig. 8.19. Where applicable, we use $S_c = 0.41$ and $G_S = 0.05$, close to the values of Eqs. (8.6) and (8.8), respectively. For simplicity, we approximate $N_A = 0$ from (8.10), and note that the results do not depend strongly on any of these parameters in the vicinity of those measured in the HAYSTAC Phase II system (Sec. 8.4). The power spectral

²⁰⁹ The essential tension is that the most obvious ways of coupling heat out of the metal rod involve also coupling microwaves out, as at low temperatures good thermal conductors tend to be good electrical conductors.




Figure C.1: (a)–(d) Signal and noise spectral densities S for the four different combinations of cold/hot rod and squeezing/no squeezing as a function of detuning ω from cavity resonance. The power spectral densities are normalized to the respective signal and noise values for the critically coupled ($\beta = 1$), on-resonant, unsqueezed case. All plots are generated using the listed coupling parameter β , obtained from the corresponding plots in Fig. C.3 so as to maximize scan rate. All plots use $\rho = \alpha = -1.0$ dB, $S_f = 0.27$, and $N_A = 0$. When squeezing (not squeezing), $G_S = 0.05$ (1). When the rod is hot (cold), $S_c = 0.41$ (S_f). Note that $\kappa_{\text{tot}} = \kappa_l(1 + \beta)$ differs for each configuration.

densities (PSDs) are ratioed to obtain the visibilities, Fig. C.2, and the scan rate R is obtained according to Eq. (5.8). By varying the coupling factor β , Eq. (8.2), we determine the scan rates for all four cases relative to the optimal ($\beta = 2$) unsqueezed, cold rod case.

The relative scan rates for the four cases are plotted in Fig. C.3, from which, together with Fig. C.1, we draw three important observations. First, a hot rod acts exactly like squeezing from the perspective of optimal coupling strategy: rather than lowering the off-resonant (reflected) noise (Fig. C.1c), it raises the on-resonant (transmitted) noise (Fig. C.1b). In both cases, overcoupling takes advantage of the lower noise PSD off resonance. For HAYSTAC Phase II parameters, squeez-




Figure C.2: (a)–(d) Normalized axion visibilities α obtained from the ratios of signal to noise spectral density curves in Fig. C.1a–d. A value of 1 corresponds to the maximum achievable visibility (critically coupled and on resonance).

	optimal β	maximum R/R_0
(a) cold rod, no squeeze	2	1
(b) hot rod, no squeeze	2.8	0.64
(c) cold rod, squeeze	4.5	1.96
(d) hot rod, squeeze	7.1	1.32

Table C.1: Theoretical parameters for the HAYSTAC Phase II haloscope with and without a hot rod and/or squeezing. Values are calculated using the parameters quoted in Fig. C.1 and the model of Fig. 8.19. (a)–(d) correspond to the panels of Figs. C.1, C.2, and C.3.

ing with a cold rod (Fig. C.3c) is optimized for scan rate at $\beta = 4.5$, while not squeezing with a hot rod (Fig. C.3b) is optimized at $\beta = 2.8$. When both squeezing and hot rod are present (Fig. C.3d), an even higher $\beta = 7.1$ is required to optimize scan rate.²¹⁰ Second, the hot rod has a little less than

²¹⁰ In practice, the insertion of the antenna necessary to achieve high coupling seemed to slightly perturb the TM_{010} -like mode of the HAYSTAC cavity.

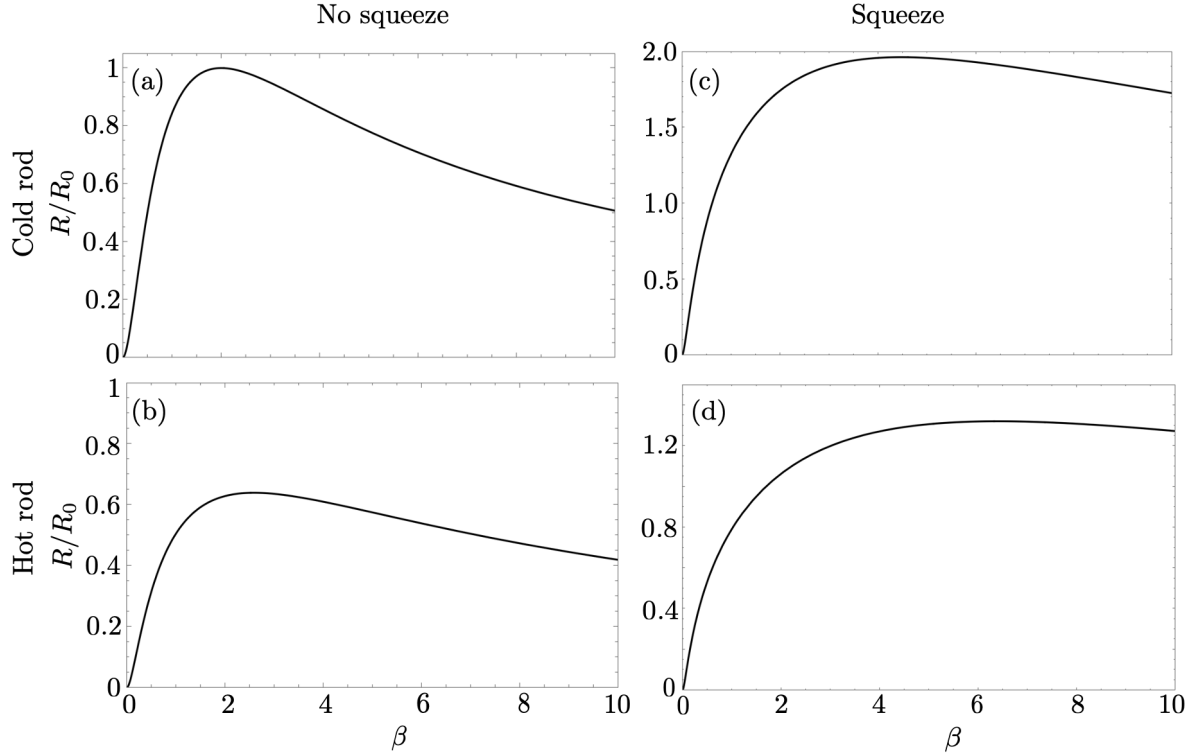


Figure C.3: (a)–(d) Scan rate R as a function of coupling parameter β normalized to R_0 , the optimal scan rate for a haloscope with neither squeezing nor hot rod, for each of the configurations shown in Fig. C.1. Higher levels of overcoupling partially compensate for a hot rod at fixed squeezing, making the scan rate enhancement due to squeezing only weakly dependent on cavity spectral density. Parameters are same as those in Figs. C.1 and C.2.

a factor-of-two impact on scan rate without squeezing present. Third, Squeezing improves both hot and cold rod configurations roughly twofold. These results are summarized in Table C.1.²¹¹ We conclude that HAYSTAC’s hot rod, while a problem nearly of the same size as the benefit conferred by squeezing, does not interfere with the ability of squeezing to improve scan rate, all else equal, by reducing truly quantum fluctuations.

²¹¹ Note that the relative scan rates R/R_0 predicted for squeezing are slightly inflated due to the $N_A = 0$ simplifying assumption used in this Appendix, which advantages squeezed over unsqueezed operation.

Appendix D

List of Abbreviations

This appendix provides an alphabetized list of acronyms and other abbreviations used in this thesis.²¹² Most abbreviations that are used across multiple chapters are redefined the first time they appear in the main text and figure captions of each chapter. Some exceptions are made where it was judged to improve readability. For example, full names for the acronyms of some experiments such as ABRACADABRA are listed only here.

Λ CDM	Λ cold dark matter (Λ is associated with dark energy)
ABRACADABRA	A Broadband/Resonant Approach to Cosmic Axion Detection with an Amplifying B-field Ring Apparatus
AC	alternating current
ACME	Advanced Cold Molecule Electron electric dipole moment
ADC	analog-to-digital converter
ADMX	Axion Dark Matter Experiment
ALP	axion-like particle
AMP	amplifier Josephson parametric amplifier
AP	axion-pertinent
AWG	arbitrary waveform generator
AWR	Applied Wave Research

²¹² Some abbreviations such as those which appear only within Figs. 2.3 and 2.4 are excepted.

BE	base electrode
BPM	Bayesian power measured
BT	Bayesian threshold
BT1	Bayesian threshold 1
BT2	Bayesian threshold 2
C4C	Center for Community
CAPP	Center for Axion and Precision Physics Research
CAPP-8TB	Center for Axion and Precision Physics Research 8 tesla B-field
Cav	cavity
CDF	cumulative distribution function
CDM	cold dark matter
CE	counter electrode
CEASEFIRE	Cavity Entanglement And Swapping Experiment For Improved Readout Efficiency
Circ	circulator
CMB	cosmic microwave background
CU	University of Colorado
DC	direct current
DEADMAU5	Dummy Experiment for Axionic Dark Matter Active Upwards of 5 GHz
DFSZ	Dine-Fischler-Srednicki-Zhitnitsky
DM	Dark Matter
EDM	electric dipole moment
eEDM	electron electric dipole moment
EM	electromagnetic
FIB	focused ion beam
FT	frequentist threshold

FWHM	full-width at half-maximum
GBP	gain-bandwidth product
HAYSTAC	Haloscope At Yale Sensitive To Axion Cold dark matter
HEMT	high-electron-mobility transistor
IF	intermediate frequency
IOT	input-output theory
IQ	in phase/quadrature
JILA	(formerly) Joint Institute for Laboratory Astrophysics
JPA	Josephson parametric amplifier
KSVZ	Kim-Shifman-Vainshtein-Zakharov
LMG	line of maximum gain
LO	local oscillator
LSSA	least-squares spectral analysis
LSW	light shining through walls
LU	least unlikely
MACHO	massive astrophysical compact halo object
MB	Maxwell-Boltzmann
ML	maximum likelihood
MOND	modified Newtonian dynamics
nEDM	neutron electric dipole moment
NIST	National Institute of Standards and Technology
OMS	one-mode squeezing
ORGAN	Oscillating Resonant Group Axion
PC	personal computer
PDF	probability density function
PMF	probability mass function
PNL	parallel nonlinear element

PQ	P eccei- Q uinn
PSD	p ower s pectral d ensity
QCD	q uantum c hromod d ynamics
QUAX- $a\gamma$	Q uaerere A xion ($a\gamma$ refers to an axion and a photon)
RBF	R ochester- B rookhaven- F ermilab
RF	r adio f requency
RMS	r oot- m ean- s quare
RWA	r otating w ave a pproximation
SA	s pectrum a nalyzer
SEM	s canning e lectron m icroscope
SG	S avitzky- G olay
SHM	s tandard h alo m odel
SHO	s imple h armonic o scillator
SN1987A	S upernova 1987A
SI	s pin- i ndependent
SIS	s uperconductor- i nsulator- s uperconductor
SMA	s ub- m iniature version A
SMP	s ub- m iniature p ush-on
SNL	s eries n onlinear e lement
SQ	s queezer Josephson parametric amplifier
SQUID	s uperconducting q uantum i nterference d evice
SSR	s queezed s tate r eceiver
SUSY	s upers s ymmetry
TE	t ransverse e lectric
TEM	t ransverse e lectrom m agnetic
TKC	t unable K err c ircuit
TM	t ransverse m agnetic

UF	University of F lorida
VNA	vector n etwork a nalyzer
VTS	variable t emperature s tage
W1	wiring 1
WIMP	weakly i nteracting m assive p article