

CAVITY NONLINEAR OPTICS WITH A COLD ATOM
ENSEMBLE ON AN ATOM CHIP:
ALL-OPTICAL LATCHING, MODULATION, AND
AMPLIFICATION

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Yeong-Dae Kwon

December 2013

© 2013 by Yeong Dae Kwon. All Rights Reserved.

Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.

<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/tq589jm9455>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Hideo Mabuchi, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Philip Bucksbaum

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

David Miller

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

I present the setup, the procedure, and the result of the experiment in which I probed nonlinear optical phenomena at very low optical power (femtojoule energy scales) in a multi-atom cavity quantum electrodynamical (cavity QED) system. I also explain the semiclassical model of this system in detail that explicitly treats heterogeneous coupling of atoms to the cavity mode, which I use to analyze the experimental data presented in this thesis. The nonlinear behaviors I demonstrated can be exploited in the design of all-optical memory, modulation, and signal amplification devices. These results highlight the interest of cavity QED systems for ultra-low power photonic signal processing as well as for fundamental studies of mesoscopic nonlinear dynamics.

Acknowledgements

First of all, I thank Hideo for all his support and guidance that made this thesis possible. I am very grateful that he prepared me a very fun project to jump in, which I have enjoyed working on throughout my graduate years and was able to finish nicely with good results. He not only provided me the resources and ideas for the project, but most importantly he advised me with the great patience when things were going slow. The work was fun but challenging, with few rewards until the end, and therefore his infinite patience was the key ingredient that allowed me to keep pushing and reach the goal at the end.

Being a graduate student in a prestigious school like Stanford is a blessing, but I think it can also be a curse. Many of us join the school with ambitious minds and can easily become obsessed with producing a lot of results at the cost of many other valuable aspects of life. In my case, thanks to Hideo and all the freedom he allowed me in doing the research, I was able to try many creative (but most of them useless) approaches for my work. Because of this, I still appreciate the pleasure of doing the experiments, and have kept the passion for continuing the path of academia. Otherwise, I would have been quickly burnt out.

Anyhow, now that I am done with my current research, I also want to thank Hideo more on a personal level aside from the work-related area. He has always been a great

leader who inspired me greatly. He treated all of us in the group with the utmost respect all the time, and throughout my graduate years, I have been deeply moved by his such attitude. I am certain that he has influenced my attitude towards others as well on some sort of a subconscious level. So, thank you.

Next, I want to thank Dr. Michael Armen. If Hideo was the father figure of the lab, Mikey was like the mother of the group, taking personal cares for everyone and making things actually work so that the lab can keep on running. The group would not be the same without Mikey – I personally think that we all should call the group not as the Mabuchi Group but as the Mabuchi-Armen group, because it is. I owe all the knowledge and technical skills that created this thesis to Dr. Armen.

I thank Joe Kerckhoff for successfully handing over the experiment setup that he inherited from Ben Lev, which formed the foundation of the current setup. I would like to say to all the former, senior members of our group, Kevin, Tony, Joe, Orion, (Lil') Mike, and Gopal that I miss you and that it was a great honor and pleasure to work with you all. Many thanks to Dmitri and Charles for always being close friends of mine, giving me the undeserved emotional support whenever needed. I was so lucky to have you as friends. Jie, Five-Stars, forever. Hardeep, you have been a great house-mate and lab-mate. Nick, Ryan, Nate, Dodd, and the rest of the lab, I liked you all very much. I wish I could have hang out with you guys more.

Mom, Dad, and my sister Yeong-Sil, I love you all, and I will do better. Sorry.

My lovely wife Yunkyong, I dedicate this thesis to you. You are the best thing that happened to me during my graduate school years, and my research results are nothing compared to your love. And you did not make me write this.

Finally, I wish to acknowledge Samsung Scholarship, which has generously funded

me throughout the four years of college and for five years during my graduate school years. Thank you very much, Mr. Lee.

Contents

Abstract	iv
Acknowledgements	v
1 Introduction	1
1.1 Overview	1
1.2 Review of optical bistability	2
1.3 Review of self-pulsing	4
2 Atom Chip	7
2.1 Work History	8
2.1.1 Atom chip with fabricated wire patterns	8
2.1.2 Atom chip with inserted micro-wires	12
2.2 Atom Chip Fabrication	13
2.2.1 Gold Coating	13
2.2.2 Wafer Saw	14
2.2.3 Wire Insertion	19
2.2.4 Attachment to the PCB	21
2.3 Advantages of the Wire Insertion Method	23

3	Atom Cloud Transportation	29
3.1	The Need for Long-Range Transportation	29
3.2	Transportation Principles	31
3.3	Some Technical Details	32
3.3.1	Transportation Speed	32
3.3.2	End of the Transportation Sequence	34
3.3.3	Fine-Tunings of the Transportation	35
3.3.4	Scattered Light	37
3.4	Atom Cloud Imaging	38
3.4.1	Fluorescence Imaging	38
3.4.2	Absorption Imaging	42
3.5	Controlling the Currents	43
4	Semiclassical Approach to Multi-Atom Cavity QED Theory	47
4.1	Maxwell-Bloch Equations for a Single Atom in a Cavity	48
4.2	Generalization to Multi-Atom Case	51
4.3	Steady State Solutions	55
4.3.1	How to Solve for the Steady State	55
4.3.2	Stability of the Steady State Solutions	57
4.3.3	Effective Atom Numbers	59
4.4	Modifications for nanophotonics applications	61
5	Optical Bistability	64
5.1	Experiment Overview	64
5.2	Input-Output Characteristics	65

5.3	All-Optical SR Latch	67
6	Limit Cycle	70
6.1	Instability Leading to a Limit Cycle	70
6.2	Sub-Critical Hopf Bifurcation	72
6.3	All-Optical DC-RF Latch	76
7	Optical Nonlinear Amplification	82
7.1	Small-Signal Amplification Using Super-Critical Hopf Bifurcation . . .	83
7.2	Measured Amplification Gain	86
7.3	Noise Analysis	90
	Bibliography	96

List of Tables

2.1	Run Sheet (1/3) : Aluminum microwire deposition steps	11
2.2	Run Sheet (2/3) : Oxide layer deposition steps	11
2.3	Run Sheet (3/3) : Gold mirror coating steps	12

List of Figures

2.1	Wire layout of the old atom chip design	9
2.2	Layers of the old atom chip design	10
2.3	Grooves on the atom chip	15
2.4	Back side of the atom chip	17
2.5	Wire insertion stage	19
2.6	Installed atom chip on a PCB	22
2.7	A photograph of a broken wire bond	25
2.8	Simulated magnetic trap	27
3.1	A schematics of the atom cloud transportation	30
3.2	Change of currents on transverse wires during the transport	31
3.3	Atom cloud compression sequence	35
3.4	Fluorescence images of an atom cloud	38
3.5	An absorption image of an atom cloud	42
3.6	Atom chip designs from Reichel Group	44
3.7	Use of the relay boxes	45
4.1	A single two-level atom and a Fabry-Perot cavity system	48
4.2	A standing wave cavity mode	54

4.3	The distribution of atom coupling strengths	56
4.4	Steady state solutions for a bistable state	58
5.1	Measured threshold and hysteretic behavior	66
5.2	Experimental demonstration of optical latching	68
6.1	Steady state solution for the system with instability	71
6.2	Demonstration of optical instability	72
6.3	Optical modulator	73
6.4	Theoretical prediction of Hopf Bifurcation points	74
6.5	Measured and predicted optical instability	75
6.6	A demonstration of all-optical DC-RF latch	77
6.7	Measured optical quadratures of a limit cycle	78
6.8	Demonstration of the sub-femtojoule dc-rf latch switch-on operation .	80
7.1	Expected power spectrum of the limit cycle	84
7.2	Power spectrum of the limit cycle at different input powers	87
7.3	Measured gain curves of an optical amplifier	88
7.4	Amplifier behaviors at different input powers	91
7.5	Decreasing amplifier output noise with larger seeding power	94

Chapter 1

Introduction

1.1 Overview

The diverse phenomena of cavity nonlinear optics [1] provide a rich basis for fundamental studies of dissipative nonlinear dynamics [2, 3] and for the design of photonic signal processing devices [4]. Recent experiments exploiting resonant atomic nonlinearities [5, 6, 7, 8, 9, 10] and/or nanophotonic cavities [11, 12, 13, 14, 15, 16, 17] have demonstrated that such phenomena can occur at very low (femtojoule–attojoule) energy scales in prototypical systems, reaching down to the quantum-physical few photon regime [18, 19] and raising intriguing prospects for corresponding ultra-low power photonic information technology [20, 21, 22]. One of the less-explored dimensions of this developing scenario is the surprising complexity of dynamical behaviors, beyond simple thresholding and bistability, that can be achieved at low energy scales in cavities incorporating two-level atoms or comparable solid-state emitters. In particular, this thesis focuses on a system composed of a Fabry-Perot cavity and laser-cooled atomic clouds, and presents a variety of nonlinear optical phenomena observed in

that system [23].

In Chapter 2 of this thesis, I explain the makings of the key experiment apparatus, atom chip, which enabled us to prepare the atom-cavity system in a way that has not been realized before and therefore allowed us to demonstrate many new phenomena in this ultra-low energy scales. Chapter 3 describes how we operated this atom chip to create such system. In Chapter 4, I outline the basic theory which can be used to quite accurately predict all the nonlinear dynamics of the system that we observed. Chapter 5 through Chapter 7 present the optical nonlinear behaviors that we measured.

1.2 Review of optical bistability

Chapter 5 of this thesis demonstrates an all-optical latching operation, using the optical bistability of a cavity QED system. We give a brief review on the optical bistability phenomenon here.

Optical bistability was demonstrated experimentally first in 1970's using a cavity filled with a thermal gas [24], and more extensively later in the 80's and 90's using atomic beams passing through the cavity [25][26]. The methods used in these early experiments were adequate for the purpose of confirming the existence of the phenomenon and verifying the match between the theory [27][28][2] and the averaged, steady-state measurements. However, because individual atoms in these experiments would be moving fast inside the cavity (transient time through the cavity waist $<1\mu s$), these were not the ideal settings to study the fast dynamics of the system reliably. The input beam power used in those early experiments were also quite large, ranging from μW to mW scale, due to the large size and the small finesse of the cavities in

use.

From the 90's, the leading experimental groups in cavity QED research turned their focus into the strong coupling regime [29]. Cavities with small mode volumes and very high finesse were used in this pursuit, and consequently the level of optical powers used in probing the nonlinearities were now finally in pW and nW scales. The first attempt to demonstrate optical bistability in this regime relied on the traditional method of using the atomic beams [5]. In this experiment, the intrinsic fluctuation in the number of atoms in the atomic beam affected the measurement results, and thus the atomic beams were deemed to be not an ideal source of the nonlinear medium for the cavity QED researches in the strong coupling regime. In order to resolve this problem, a way of slowly transporting laser-cooled atoms into the cavity mode were required. The two major techniques were developed, the one using the magnetic trapping on atom chips [30] and the one using 1-D optical lattice trap with strong, far red-detuned light [7][31], both of which have been very successful in leading the cavity QED research in the strong coupling regime.

Looking back, when the focus of the cavity QED research community shifted from the weak coupling regime with low finesse cavities ($n_0, N_0 \gg 10$, where n_0, N_0 are the critical photon number and the critical atom number [19]) to the strong coupling regime with high finesse cavities ($n_0, N_0 < 1$), the intermediate regime ($1 < n_0, N_0 < 10$) was left unexplored during the transition. And it is precisely this intermediate regime that we probed and report in this thesis. With the increasing interest in the nanophotonics community in the development of ultra-low power all-optical devices [20, 21, 22], the importance of better-understanding of this intermediate regime would need to be highlighted, because this regime will be the operating

domain of all-optical signal processing devices as the progress is made towards the theoretically lowest switching energy. Once you go below this limit, you inevitably fall into the strong-coupling regime, where the quantum noise will compromise the stability of the classical operation of the device [32][19].

Also, unlike other cavity QED experiments in the past that only focused on observing the hysteresis with a slowly sweeping the input power, we attempted for the first time to measure the minimum switching energy (minimum temporal width of the switching pulse) of the bistable system by applying random sequences of input pulses, as measurements of such switching energy is one of the key interests in the nanophotonics device community [33][14].

It should also be noted that demonstrations of optical bistability has been published at even lower input power scales in the past [5][7] compared to that of our experiment, but these results were taken from the cavity QED systems in the strong coupling regime. Therefore, as pointed out earlier, quantum fluctuations (and other technical noises) led to incomplete matches between the theory and the measurement in these experiments. Moreover, no attempts to demonstrate a real-time latching operation over an extended time were made at the time, not to mention the measurement of the minimum switching energy itself, presumably because a reliable switching operation with a reasonably small energy was not possible in these strong coupling regimes.

1.3 Review of self-pulsing

Chapter 6 of this thesis is about a limit cycle phenomenon observed in our system. Instability and chaos in the cavity QED system has been the focus of many research

groups in 1980's, and such instability was experimentally demonstrated using various methods, such as an electrical delay in the feedback loop [34], a Q-switched laser [35], magnetically induced Larmor precession of the atoms [36], or multi-mode cavity [37]. Later, it was predicted that the instability can arise intrinsically (external drive being the cw-input only) in a single mode cavity QED system using an appropriate cavity and laser detunings [38], and this self-pulsing behavior was soon experimentally demonstrated using atomic beams crossing the cavity mode [39][40]. The input beam power used in these experiments were in mW scale.

Unlike the phenomenon of optical bistability, however, the research on this self-pulsing phenomenon has been discontinued since then, and has never been demonstrated in any other settings of the cavity QED system until now. This is partly because the cooperativity (or the number of coupled atoms) you need for the self-pulsing is one or two orders of magnitude higher than what is required for observing optical bistability, and thus it is relatively difficult to achieve. The result of our experiment presented here, therefore, represents the observation of self-pulsing behavior in a cavity QED system with a lowest optical input power level ($\sim 1\mu\text{W}$) so far, in addition to the fact that it was the first measurement of the required energy to induce switching between the dc and the rf state. We have also verified the good match between the theory and the limit cycle in a standing-wave cavity mode for the first time, as the previous experiments have only verified such match in the ring-cavity setting [40].

Interestingly, the self-pulsing behaviors exhibited in a nanophotonic devices have now started to draw attentions of the nanophotonics community. The self-oscillations

in these photonic devices are induced by various methods, such as coupling two micro-cavities [41], using two-photon absorptions [42], or taking advantage of the band structure of graphene [43]. The intrinsic method used in our experiment has not been adopted to the nanophotonic devices yet, even though the bistability phenomenon in the cavity QED has been successfully adopted and widely studied in the nanophotonics settings by now.

Chapter 2

Atom Chip

Atom chip [44, 45, 46, 47] is a technique developed in the 1990's to trap and transport neutral atoms at the near-surface of a chip. Unlike the 'optical' trap/transportation technique [7, 31], the atom chip traps atoms 'magnetically', using the magnetic field gradient induced by electrical currents underneath the surface. This leads to many advantages over other methods when trying to couple atoms to various optical devices on the chip's surface. First, you do not need to worry about placing the experimental components in a way that does not block the trapping/transporting beams. You can even transport atoms in a curved trajectory [48] giving you more room to play with when arranging optical components. Second, and more importantly, atom chip technology can be naturally merged with various nanophotonic devices as both devices are based on a fabricated planar chip geometry.

Coupling atom-like optical components, such as quantum dots and NV-centers to nanophotonic devices is already an active focus of many research groups [12, 13, 15, 14]. Because the 'real atoms' have superior coupling strength and longer coherence time compared to these 'atom-like' particles, we can probe many interesting physics

using atom chips before the field of nanophotonics gets matured and similar phenomena are observed using nanophotonic devices, as is demonstrated in this thesis.

In this thesis, I will not go into the details of how the atom chip works. Refer to Ben Lev's thesis [49] or other references mentioned above to better understand the physics behind it. This chapter is mainly dedicated to explaining the atom chip fabrication method I used for my experiment, which was a little bit unconventional.

2.1 Work History

2.1.1 Atom chip with fabricated wire patterns

Back in the old days when Mabuchi Group was still at Caltech, one of the former graduate students of the group, Ben Lev, had been leading the atom chip experiments. He graduated as an expert in atom chip fabrication in 2006 [49], but at the time the group had no one else to continue from where he left off. Most of his skills and knowledge on the atom chip fabrication was lost by then¹.

After Ben left, the atom chip experiment setup was kept alive thanks to Joe Kerckhoff, another former graduate student, but the atom chips themselves were now no longer house-made. A group at Sandia National Lab, through the collaboration with our group, shipped us some of their prototype atom chips that they were working on so that we can use our setup to test them. I joined the Mabuchi Group in summer of 2007, during the transition period when our group was moving from Caltech to Stanford. And near the end of my first graduate years (2008 summer), I inherited the atom chip project. Beside learning and improving on the existing optical setups used

¹I inherited a drawer full of Ben's fabrication related materials, such as sapphire chip, lithography masks, etc. I doubt that they would ever be used again.

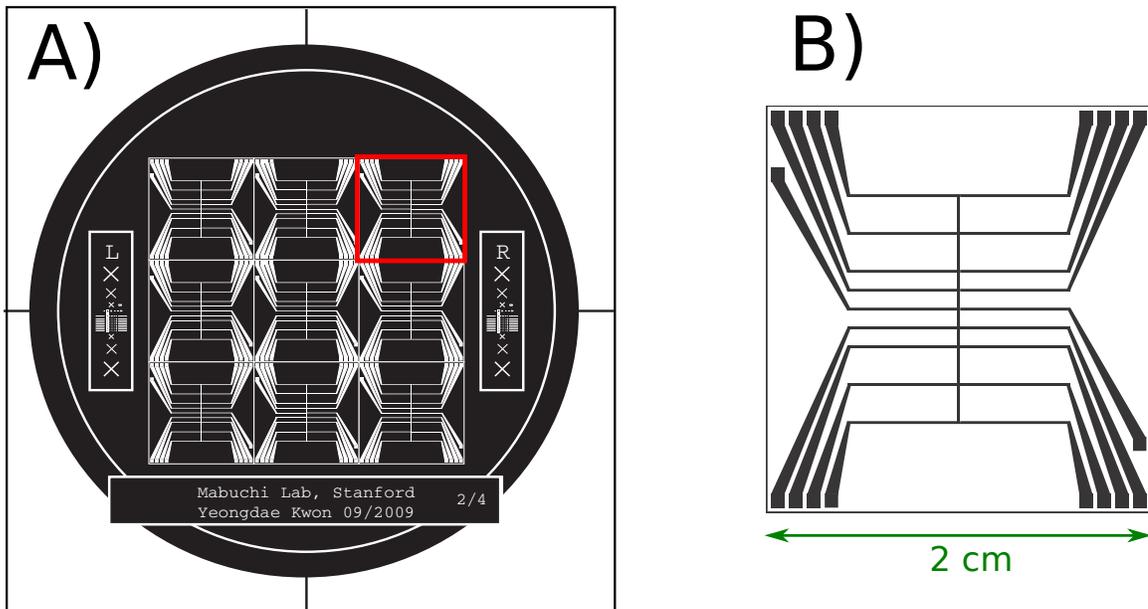


Figure 2.1: **A)** One of the photolithography masks I used for the atom chip fabrication. Because the feature size I needed was large (100 micron for micro-wires), it was okay to simply order these designs to be printed on a transparency film, and use it for UV exposure during the photolithography. One wafer contains 9 copies of the atom chip. One copy is highlighted by a red square. **B)** Micro-wire pattern used for my first atom chip.

for atom chip experiments, one of my early tasks was to revive the group tradition on producing house-made atom chips, this time using the advanced fabrication facilities at Stanford University such as Stanford Nanofabrication Facility (SNF) Labs.

The first atoms chip I made was based on the Sandia Lab design. The wire pattern embedded in the chip is shown in Fig. 2.1B. This was a prototype device and coupling atoms to optical resonators was not yet attempted. Rather, its immediate goal was to test MOT making, atom trapping, and possibly a somewhat limited version of atom transportation that relied on diffusion of atoms in a 2-D trap [49].

Figure 2.2 shows the basic structure of the atom chip. On the polished side of

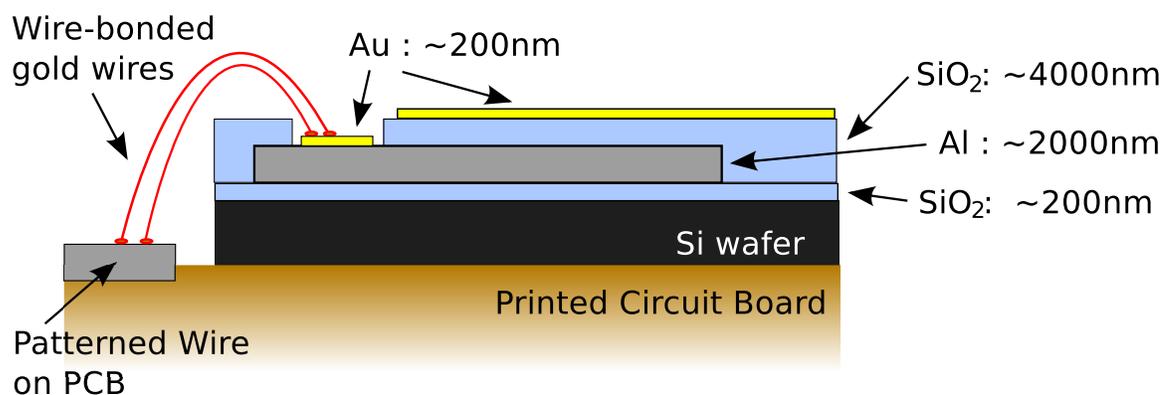


Figure 2.2: Structure of the early atom chip with fabricated wires.

the silicon wafer, 200 nm oxide was grown as an insulating layer and then 2 micron thick aluminum wire patterns were created on top of this layer. The wire patterns were then covered by 4 micron thick oxide layer, except for places where electrical connections were made. Finally a thin gold layer was deposited to make a mirror (reflective) surface. After the fabrication procedure, the wafer was cut into nine square pieces (see Fig. 2.1A.) and one of these square chips was glued onto a printed circuit board (PCB) using a vacuum compatible epoxy (EPO-TEX 353ND). Lastly, electrical connections were made by wire bonding.

Major fabrication steps I followed in making the atom chip are listed in Tables 2.1-2.3. Even though the idea behind the structure of the atom chip is a very simple one, namely, patterned wires underneath a flat mirror surface, the actual fabrication procedure to accomplish this was considerably more complicated than what it seemed to me at first. As you can see, the hours needed to go through all those steps on the list are quite significant, not to mention the extra training and/or shadowing hours that are mandatory for the use of each equipment.

Equipment	Time	Summary
n/a	n/a	Scribe on the back of the wafer to label.
wbnonmetal	0.5 hr	'Phirana' clean of a new silicon wafer in a wet bench.
wbdiff	1 hr	Standard prediffusion clean in a wet bench.
tylan1	2.5 hr	Grow 200nm oxidation layer into the silicon using a furnace.
nanospec	n/a	Confirm oxidation and measure the oxide thickness.
Gryphon	1.5 hr	Deposit 2 micron of Aluminum layer.
P2	n/a	Measure Al thickness.
litho tools	4 hr	Photolithography to mask the wire pattern using photoresist.
wbmetal	1.5 hr	Wet etch Aluminum, and then strip photoresist.

Table 2.1: Run Sheet (1/3) : Aluminum microwire deposition steps

Equipment	Time	Summary
wbmetal	0.5 hr	Standard metal clean in a wet bench, before using the furnace.
tylanBPGS	8 hr	4 micron Silicon dioxide deposition, using low-temperature, chemical-vapor-deposition furnace. Dummy wafers used in the furnace must be cleaned afterwards in the wet bench (wbsilicide and wbdiff).
nanospec	n/a	Measure the oxide thickness.
P2	n/a	Measure the height of the oxide bump, created by the existing thick wire patterns underneath the deposited oxide.
litho tools	4 hr	Photolithography to negatively mask a little thinner wire pattern using photoresist.
wbmetal	2 hr	Etch out oxide bumps on the wafer surface, and strip photoresist.
CMP	3 hr	Planarize the surface. (Chemical Mechanical Polishing/Planarization.)
wbsilicide	3 hr	KOH decontamination after CMP

Table 2.2: Run Sheet (2/3) : Oxide layer deposition steps

Equipment	Time	Summary
litho tools	4 hr	Photolithography to negatively mask electrical connection pads at the end of the micro-wires.
wbmetal	2.5 hr	Etch out oxide layer to expose the Al pads (using etchant called 'Pad Etch'), and strip photoresist.
P2	n/a	Measure steps created by the etching and confirm that Al is exposed.
n/a	n/a	Confirm electrical connections using a multimeter.
litho tools	4 hr	Photoresist mask for gold lift-off.
innotec	2.5 hr	Deposit 10nm Ti (adhesion layer) and 200nm Au.
wbsolvent	0.5 hr	Use Acetone to lift-off gold layer.
litho tools	2 hr	Coat wafer surface with photoresist to protect it from the debris in the next step.
wafer saw	2 hr	Cut wafer.
wbsolvent	0.5 hr	Use Acetone to wipe out photoresist protection layer.

Table 2.3: Run Sheet (3/3) : Gold mirror coating steps

2.1.2 Atom chip with inserted micro-wires

Other than the method described in the previous section, there are various of different methods that can be used for fabricating atom chips. Instead of using an oxide deposition to create an insulation between the wires and the mirror surface, one can use polyimide to easily make a corrugation-free surface without the planarizing process, because you can simply spin it on the wafer to create a flat surface while it is still viscous and then harden it afterwards [50]. One can also prepare a thin mirror surface and a separate substrate surface with patterned wires and epoxy them together [51]. Rather than using separate layers for the wires and the mirror surface, one can even make an atom chip with a single gold layer; you can simply gold-coat the wafer surface to make a mirror and use that same gold layer also to create wires, with gaps to outline the wires so that they are electrically insulated from the rest of

the gold surface [52].

After I made the first batch of atom chips following the procedures as explained in the previous section (late 2009), I switched to a yet simpler way to produce the atom chips. The method I developed is somewhat unconventional. Rather than fabricating the micro-wires directly on the chip surface, I chose to use a spool of copper wire purchased from a manufacturer. I cut the wires into the proper lengths, and manually inserted them into the grooves made on the back side of the gold-coated chip. Once the wires in position, I use epoxy to keep them in places. A similar (but not quite identical) approach to make an atom chip has been used in another paper [53]. All the experiment results that will be presented in the later part of this thesis have been obtained using the atom chip made in this fashion.

2.2 Atom Chip Fabrication

In this section, I will describe in full detail how I produced the atom chip using the wire insertion method, briefly outlined in the previous section.

2.2.1 Gold Coating

The first step of making an the atom chip by the wire insertion method is to prepare a wafer with a gold mirror coating on its front side. After I cleaned a brand new silicon wafer², I used an e-beam metal evaporation system (Innotec) to deposit 10 nm Ti (adhesion layer) and 150 nm Au. The atoms chips used in my experiment were made from this first batch of gold-coated wafers.

²Dumped in a Phiranha solution (90% sulfuric acid (H_2SO_4) and 10% hydrogen peroxide (H_2O_2)) for 20 minutes. I was simply following the protocol of the SNF Lab, but I wonder if this is really necessary in my case because I was making a mere gold mirror.

I also made a second batch of gold-coated wafers at some point, this time finishing it with an extra 200 nm silicon dioxide (SiO_2) layer added on top of the gold layer. The reason I tried this second batch was because I noticed that the bare gold layer of the first batch got easily scratched, especially when I attempted to clean its surface by wiping it with a lens paper soaked with acetone. Pure gold is a soft metal, after all. So I made this second batch, thinking that I would test the surface and if satisfactory I could use it for the next atom chip fabrication. The surface with this protective SiO_2 layer turned out to be very solid and almost completely scratch-proofed. I was very much satisfied, although I never got the chance to make a new atom chip using this second batch.

One other comment I would like to make is that since gold-coating is such a simple procedure, rather than getting the training for the metal deposition equipment, one could probably easily ask someone else with the required training to do the job for them. Making an order to some outside company for such a simple job can't be so costly either.

2.2.2 Wafer Saw

The next step is to use a wafer saw to make grooves to insert wires and also to cut out the outline. Normally when people use wafer saw to cut out a die from a wafer, a blue adhesive film is applied on the backside of the wafer to hold it during the dicing process. However, in my case, I needed the grooves (dicing streets) on the back side of the wafer, and therefore my wafer needed to be placed on the wafer saw face down, with the adhesive film glued to the front side (mirror coated side). To prevent the adhesives film leaving a residue or damaging the mirror surface, I applied

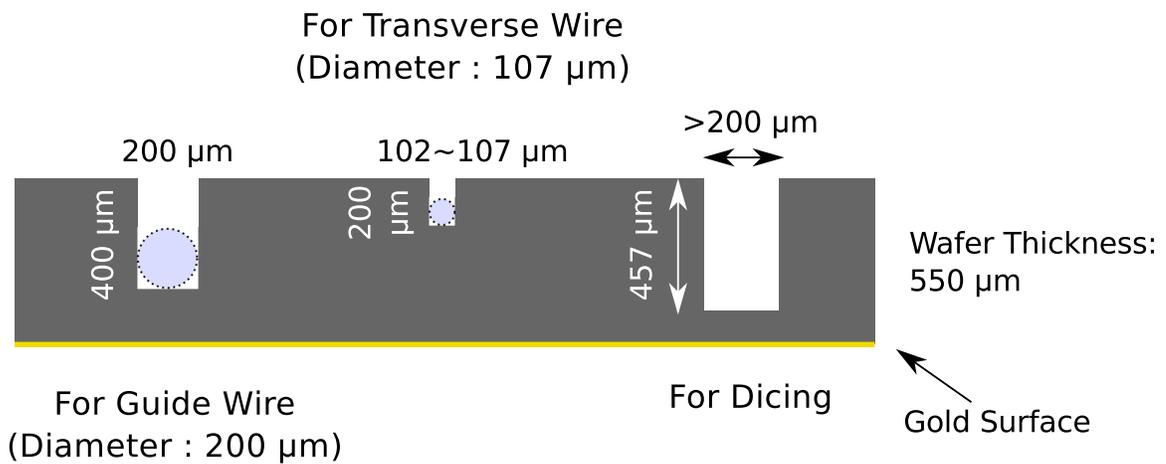


Figure 2.3: Three types of the grooves that I make on the back side of the wafer by the wafer saw.

a thick overcoating photoresist on the gold surface and hard-baked it before attaching it to the film. After dicing and the wire-insertion, the film was peeled off, and the photoresist was cleaned using acetone.

Types of Dicing Streets

Figure 2.3 shows three different types of dicing streets I used in the latest atom chip design. The first type is the large groove meant for the guide wire with diameter of 200 μm, and the second type is the smaller groove for transverse wire with diameter of 107 μm (copper core diameter 100 μm, and an added thickness of 7 μm due to the insulation). As is explained in 2.3, it is advantageous to bring the wires as close to the surface as possible. Therefore, the depth of the guide wire is chosen to be reasonably deep (400 μm), but not so deep that there's a danger of breaking the wafer along the cut. The depth of the transverse wire is chosen, in turn, to 200 μm to ensure that it can run across the guide wire without the physical interference. The widths of

these two types of the grooves were meant to suit the diameter of the corresponding micro-wires. I found that a tight fit makes wire-insertion procedure a little easier because there's less movement of the wires after the insertion. This trick of squeezing the micro-wire into a tight groove works best for transverse wires but not so desirable for guide wires, because if we try to squeeze the guide wire into the wide and deep groove it could break the chip.

One cut operation by the wafer saw will create a dicing street whose width is determined by the thickness of the blade. In my case, I used the standard blade used in the SNF Lab, Semitec F3530, whose thickness is $80\ \mu\text{m}$. In order to create dicing streets whose widths are larger than $80\ \mu\text{m}$, I made double- or triple-cuts at the same location with some appropriate displacements of the cut position³. It means that for the second and third cut, one side of the blade will be in an open space, while the other side of the blade will dig into the side-wall and expand the groove width. In a sense, I am using the wafer saw as a milling machine. Because this technique induces an unbalanced stress for the blade, it tends to shorten the blade's lifetime. I found that the actually cut width becomes larger over time ($80\ \mu\text{m} \rightarrow 100\ \mu\text{m}$ or more). But other than this side effect, which is a minor one since we do not need an accurate cut width, this double- and triple-cut technique worked quite well for my purpose. There are also companies that sell wafer saw blades that have larger than $80\ \mu\text{m}$ thickness, if double/triple cut should be avoided.

The third type of the cut I make using the wafer saw is the 'ordinary' dicing street, a deep cut ($457\ \mu\text{m}$) that will form the edge of the atom chip (Figure 2.3, third from the left). The cut depth was chosen to match the maximum depth specified for

³The exact location of the dicing street can be controlled by the wafer saw down to better than 1 micron.

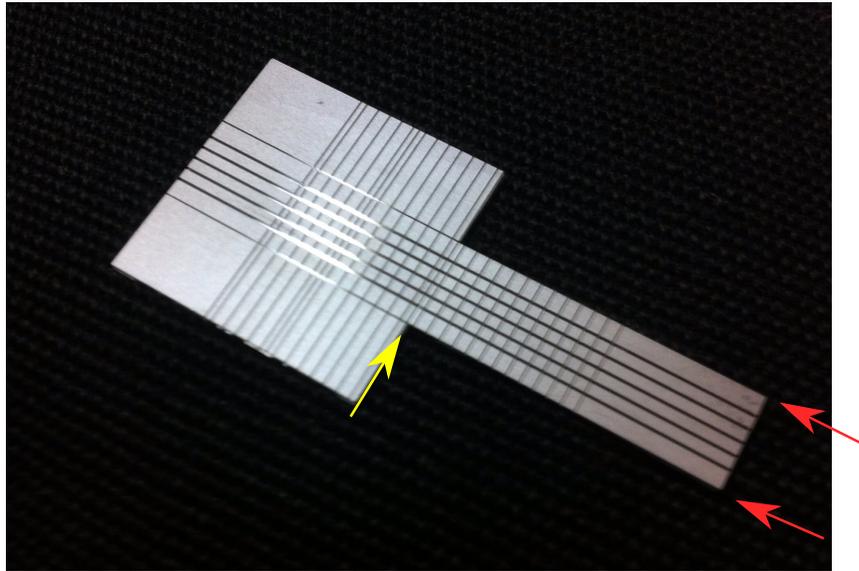


Figure 2.4: Back side of the atom chip, before inserting micro-wires. Five parallel grooves that run horizontally (in the long dimension) are the dicing streets meant for inserting guiding wires, and the thinner grooves that run perpendicular to those are for the transverse wires. During the experiment, the atom cloud is transported along the center guid wire, on the other side (mirror surface) of the chip. The separation between the five horizontal grooves is one millimeter. Yellow and red arrows are explained in the text.

Semitec F3530 blade I used. After I am done with all the cuts, I rinse the wafer in running water and dry it, and then I apply a gentle pressure to the wafer to split it along this line. Normally, a single dicing street, regardless of its width, would be sufficient for snapping the wafer along the line this way. But in my case there are many more dicing streets (for wire insertion) already made on the wafer (see Figure 2.4) and the pressure to the wafer can easily crack it at the wrong line, that is, not along the real dicing lines but along these wire-insertion lines. Therefore, I made the dicing street wider than the wire insertion grooves ($> 200\mu\text{m}$) to make it the weakest point. Also, it was important that I didn't rely on my fingers to snap the atom chip

wafers after the wafer saw, which is a common practice in the typical situation. I laid down the wafer on a flat surface with an edge (like at an edge of a table or at an edge of a thick metal block) and used another piece of wafer with a straight edge to apply pressure at the exact location and allow only the wanted line to break.

Unavoidable Cracks

When creating an atom chip with a unconventional shape that is not rectangular, it sometimes happens that making cracks on the chip in a few places is unavoidable. In the case of the latest chip design (shown in Figure 2.4), the chip had to be snapped along the line marked by the yellow arrow. Only then, I could cut the wafer along the lines indicated by the red arrows to complete the shape I needed. Fortunately, the cracks on the atom chip are not as bad as they sound, because they can be easily fixed. The blue sticky film is still attached on the front side of the wafer when we snap it, and therefore it will holds all pieces together in their initial places even though they are now physically severed from one another. After I peel all the unnecessary wafer pieces off the film, leaving only the ones that compose the atom chip, I can apply epoxy to flow through the (cracked) dicing streets between the pieces and cure. This will glue all the separated pieces into one again. This leaves a fine line on the gold mirror surface, nevertheless, which will affect the uniformity of the reflected MOT beam and therefore affect the forming of the MOT. The crack line along the yellow arrow in the figure, however, is outside the MOT forming area. As for the transportation of an atom cloud, this crack line does not matter, since the transportation is done purely magnetically.

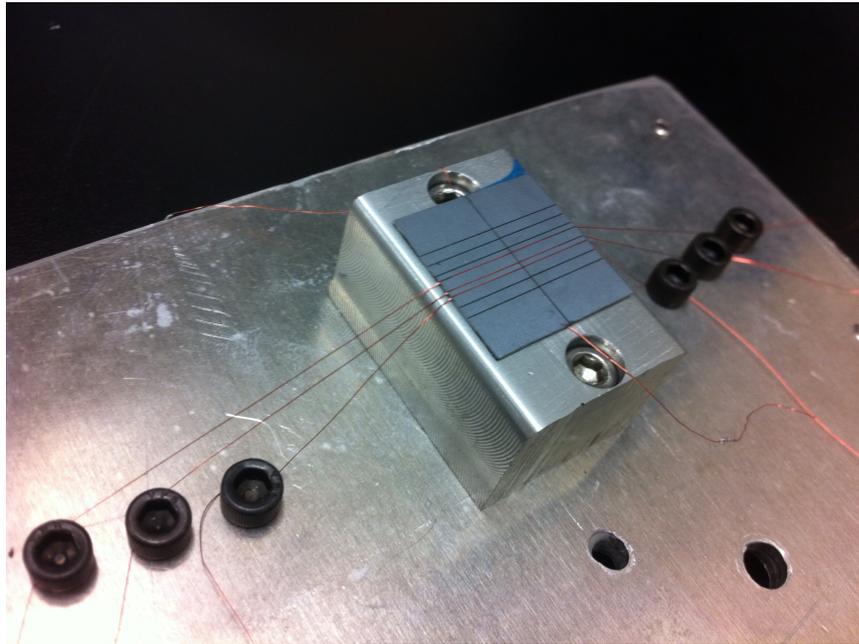


Figure 2.5: An example that shows how the micro-wires insertion works on a house-made stage. In this picture, three transverse wires are being installed in an atom chip. (The chip is of one of the old designs.) The wires are stretched tight and held by the screws, and it is now ready for a few drops of epoxy. With the epoxy filling the space between the wires and the grooves, the whole stage is loaded into an oven for the curing process. You can also see the guide wire (a single, vertical wire in the picture) that has been installed at the deeper groove before the insertion of the transverse wires at the shallower ones. Because it is already epoxied, it no longer needs screws to hold it down.

2.2.3 Wire Insertion

Once the chip is cut out from the wafer, the next step is to integrate the micro-wires into the chip. I would insert the wires into the bottom of the grooves and lock them into the place by epoxy (EPO-TEX 353ND). The main difficulty during this process is keeping the wires in position while the epoxy around them hardens. For this purpose, I assembled a stage that will hold the wires and keep them under tension. A picture of this stage is displayed in Figure 2.5. Screws were used to bolt down both ends of

the micro-wires while its center was inside the chip. It was important to have the chip raised up from the wire-holding plane so that the tension would keep the wires pushed down, entire segment inside the chip touching the bottom of the grooves.

The stage was also made to be portable and high-temperature tolerant. Once the wires are bolted down, epoxy is applied and then the whole stage would be loaded into an oven to heat and cure the epoxy. EPO-TEX 353ND has the cure time of 1 minute at 150°C or 30 minutes at 80°C, and over many trials I settled to a low temperature cure as it seemed to decrease the chance of cracking the atom chip during the baking process. If not careful, the tension of the wire can break the atom chip in halves along the groove. A high temperature and/or a rapid change in the temperature seemed to facilitate such disaster⁴. For an atom chip design that has a few dozen wires to be inserted, glued, and cured, the whole insertion process can take many days unless you can process multiple wires at a time (as shown in Figure 2.5).

For the micro-wires to be inserted into the grooves of the chip, I used magnet wires bought from MWS Wire Industries. Magnet wires are meant for a tight winding around something like a conducting rod to make an electromagnet, and therefore they come with a small diameter and with a very thin insulation layer on the outside to maximize the winding density. Fortunately, these are also the characteristics that we need for the atom chip fabrication. The 38 AWG wire that I used for transverse wires has a copper core with a diameter of ~ 101 micron. The vendor sells wires as thin as 60 AWG, which has ~ 8 micron diameter. Thankfully, the vendor provided us a few meters of any types of wires as a sample upon request, with which I could test the wire properties or practice inserting them into the chip.

⁴As I mentioned before, though, crack lines on the atom chip surface does not mean that the chip is unusable. But I still tried to minimize the cracks.

As for the insulation coating, the vendor sold magnet wires either as a bare copper or with a few different types of coating. My latest atom chip design used insulated wire to make the transverse wires in order to prevent a short (but bare copper for the guide wire since there was no such danger for the guide wires), and it had “Single polyurethane Nylon” coating. I wasn’t very careful when choosing the insulation coating type, but luckily this coating did not cause any problem under the ultra-high vacuum. There is no significant outgassing from the coatings in general, as it would have prevented me from pumping the chamber down to the UHV (10^{-10} Torr range). I also tested and saw that there was a increase in the background pressure if I run too much current for too long and heat up the wire, but during the normal operations no such rise in the pressure was detected.

2.2.4 Attachment to the PCB

After all the micro-wires are properly installed and bonded to the chip, the atom chip can be now mounted on a PCB. First, the atom chip itself is glued to the PCB to secure its position. Then the magnet wires that now stick out from the sides of the chip are cut to the proper length, and then the insulation coating on their ends is scratched off by a razor to expose the bare copper. This exposed copper is soldered directly onto the metal pad on the PCB by the fluxless solder using an ultrasonic soldering iron. Soldering with flux is easier, but it is avoided in order to reduce the outgassing when it is placed inside the UHV vacuum chamber. This completes the atom chip fabrication process using the wire insertion method.

The wire layout of the PCB was designed in the lab using a electronics design software called Altium Design Explorer (DXP). The latest atom chip design utilizes

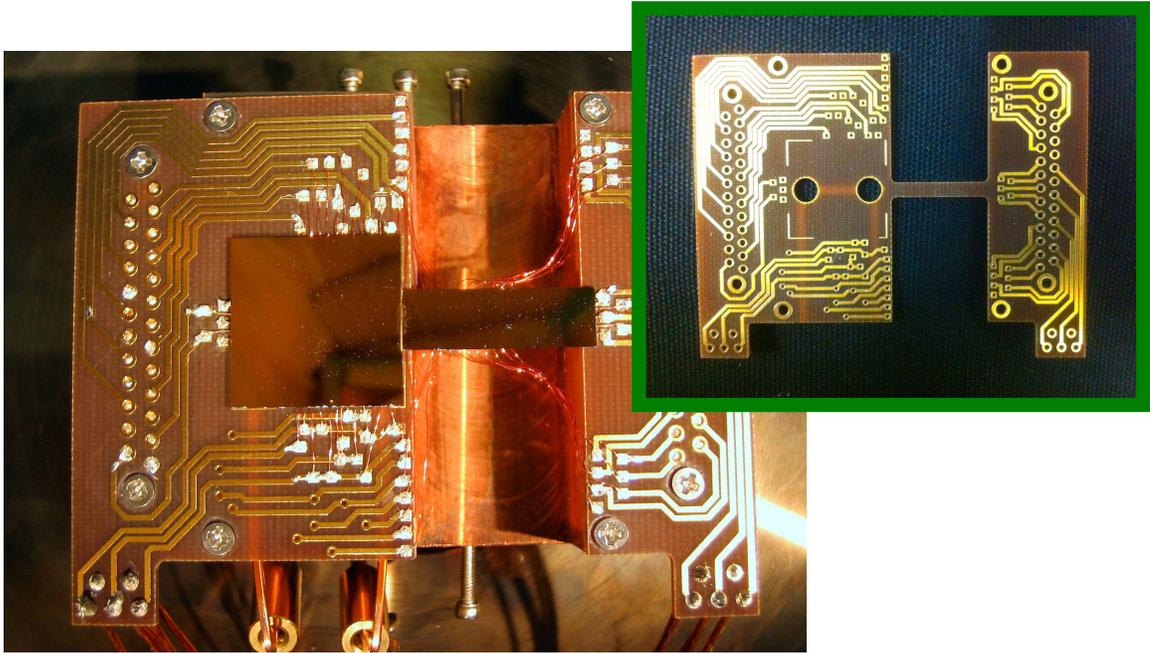


Figure 2.6: Installed atom chip on a PCB. The photo was taken after it was mounted on the copper block that goes inside the vacuum chamber. Inset on the upper right corner shows the bare PCB before the atom chip was installed.

two layers (Top and Bottom) to arrange all the necessary traces that will connect the micro-wires (5 guide wires and 27 transverse wires) to the connector pins to be installed on the PCB. (See Figure 2.6). Once the design on the computer is done, Altium-DXP can generate output files that we can send to a PCB shop to have it printed. We make the special request so that the PCB is to be manufactured using ‘polyimide material’, which is UHV compatible and able to withstand the high temperature ($150^{\circ} \sim 200^{\circ}\text{C}$) while the chamber is baked. Two 25-pin male D-connectors are installed on the PCB, whose parts are purchased from Accu-Glass Products (part number 100420). They are made with polyether-ether-ketone (PEEK) material, which is also UHV compatible and heat-resistant.

2.3 Advantages of the Wire Insertion Method

Simpler and Faster

The most obvious and significant advantage of the wire-insertion method over the wire-fabrication method is its simplicity. The time required to make one atom chip is greatly reduced. Given that all the necessary parts are prepared, an atom chip can be made within a few days using the wire-insertion method, while it could take weeks with the wire-fabrication method. If the wire-insertion method is used, not even one photolithography is required, because the relative positions of the grooves (and hence the positions of micro-wires) can be enforced reasonably accurately by the computer-controlled wafer saw.

Moreover, the gold deposition process onto a bare silicon wafer, the only fabrication processes that now need to be done inside the clean room, does not need to be repeated every time a new atom chip with different design is attempted. One can produce a dozen of gold-coated wafers all at once and stock them for future change of atom chip design.

High Current Tolerance

Another major advantage of the wire-insertion method is the increased durability of the atom chip. When you fabricate the micro-wires by depositing metal on the chip and patterning it, there is a limit on how thick a metal layer you can grow (without someone getting angry at you for abusing the machine) that sets the maximum height of the micro-wires for the atom chip. In the SNF lab, I needed to ask for special permission to sputter-deposit more than 2 microns thick aluminum. The equipments

in the fabrication facility are generally meant for nano-scale operations, not micron-scales. Given a fixed amount of current that you need to run, a micro-wire's small cross sectional area directly translates into an increased electrical resistance (i.e. increased heat generation), and slower heat dissipation rate. Micro-wires do get blown out very easily if the amount of electrical currents is not regulated very carefully.

Generally speaking, the dimensions of the micro-wires are best kept small, because they determine the accuracy of the trap position. I believe that the ideal height of the micro-wires should be on the same order of what its width is, and that is the case when you use the magnet wires. (They have a circular cross section⁵.) When I tested the magnet wires with the diameter of 100 microns (38 AWG) inside the vacuum chamber, I noticed that I can safely run tens of amperes of current for a few milliseconds, or a little more than one ampere for almost indefinitely, without noticing any change in the wire resistance or the background pressure inside the chamber. Insulation coating around the wires can burn out when the wire temperature goes too high and leads to increased pressure. In the real experiments, I would sometimes run 2 or 3 amperes through these wires for ~ 100 milliseconds without too much concern.

No Wire Bonding

High durability of the wire-insertion method also comes from the elimination of the wire-bonding process. (See Figure 2.2 and 2.7.) If the micro-wires are fabricated on the atom chip, the electric connections between these fabricated wires and the wires on the PCB must be made by wire-bonding. However, the bonded wires are very

⁵Magnet wires with square cross section also exist. If I could have used them for atom chips, they would make wire-insertion method even easier because they would fit inside the groove even better. However, the options for square-shaped wires were only for those with larger width.

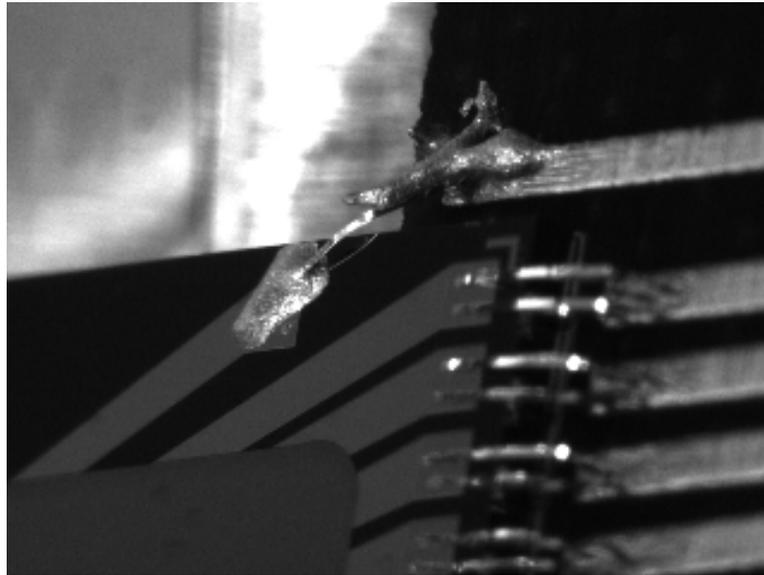


Figure 2.7: A photograph of an atom chip with the wire-bonded gold wires connecting the chip to the PCB. This chip was made in the Sandia Lab and shipped to us. Before I tested the chip inside the vacuum chamber, I wanted to clean the chip for any possible dirt on the surface, so I ran an ultrasonic clean. This turned out to be a terrible idea, because the mechanical vibrations from the cleaning process destroyed the wire-bondings in a several places. I had to fix the broken wires by applying conductive epoxy (EPO-TEK H20E), as shown on the top of the picture, in focus. Below the fixed wire, you can see the undamaged gold wires connecting the chip and the PCB.

fragile and are very vulnerable to external stress. The connections made by wire-bonding also break easily when too much current was applied to the micro-wires by mistake. On the other hand, when the wire-insertion method is used no such wire bonding is required. The micro-wires are soldered directly onto the PCB wires and these connections are very strong both physically and thermally.

Separation of the Guide Wire and the Transverse Wires

When you fabricate all of your micro-wires on one layer, as in Fig. 2.1B for example, complications arise because all the wires are interconnected. When multiple power

supplies are attached to such an atom chip, controlling the currents on micro-wires with multiple input and output ports becomes a very hard problem. Because the current generated by one machine can sink into the ground of another, the power supplies can no longer regulate its outputs properly. One may try to solve this problem by creating a virtual common ground for all the generators on floating-ground settings, and carefully adjusting their input/output voltages according to the resistance of each section of the wires. I have not attempted this, but I have no doubt that this will be an extremely difficult task.

This problem can be avoided if one uses more than one layer to fabricate the micro-wires and avoid crossings, of course. But adding an extra layer increases the workload of the fabrication process significantly.

Wire insertion technique, on the other hand, naturally avoids this interconnection problem. By using the magnet wires with the insulation coating, the electrical isolation between the guide and the transverse wires is easily achieved. Each wire can be assigned to a different power supply and the current through the wire can be controlled independently, and reliably. Figure 2.8 shows a simulated atom trap formed by currents through three individual wires and a proper bias magnetic field.

Trade-offs

Generally speaking, an atom chip with finer micro-wires means more precise control on the trapped atom position, because there is less ambiguity on the exact location of the magnetic trap. It also means that the atoms can be brought closer to the center

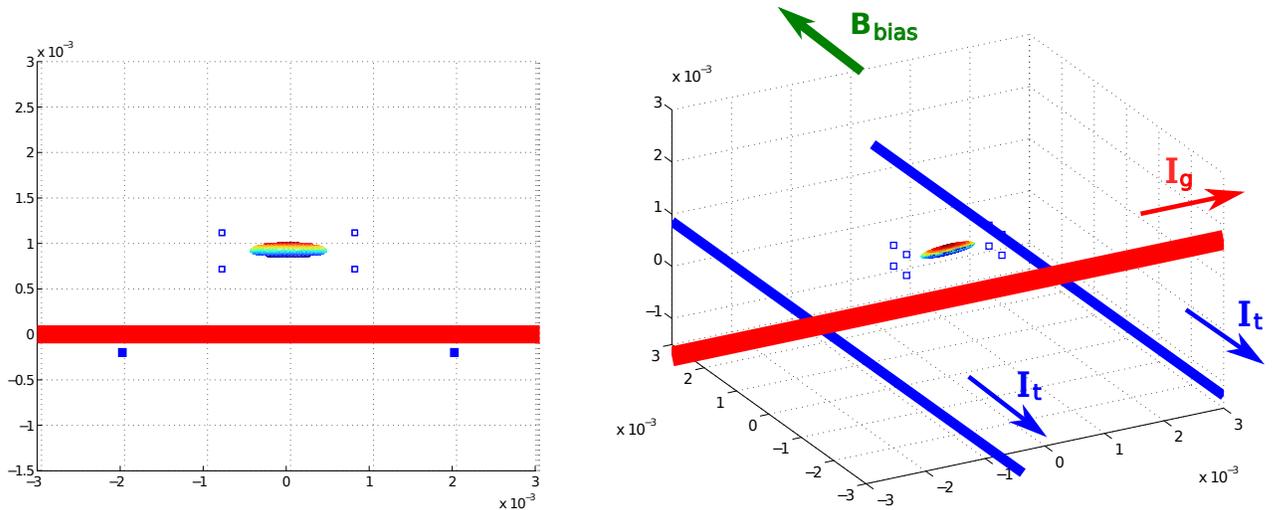


Figure 2.8: Simulation of the magnetic trap, formed by the magnetic fields around three independently running wires and a strong bias field. The elliptical contour on top of the guide wire (red) and between the pair of transverse wires (blue) is the calculated atomic trap boundary. In the simulation, I ran 4 amperes through the guide wire and 2 amperes through the transverse wires. Bias magnetic field was 9 gauss. The trap boundary is calculated for the atoms at $10\mu\text{K}$.

of the wire. The magnetic field around the wire is given by

$$B = \frac{\mu_0 I}{2\pi r}, \quad (2.1)$$

and therefore the closer you get to the center of the wire, the stronger the field gradient becomes. This way, you can make a tighter trap, increasing the atom cloud density and its lifetime.

Although I used relatively thin wires ($100 - 200\mu\text{m}$) in my atom chip with the wire-insertion method, these wire diameters are orders of magnitude larger than what is easily achievable by the wire-fabrication method. If the goal of the experiment was

to couple a single or a few atoms with a nanophotonic resonators, the wire insertion method would not be the ideal choice to make. However, the main function of the atom chip in my experiments was to transport a large number of atoms over a long distance and couple them to a resonator that has a relatively large mode volume. The atom chip made by the wire-insertion method worked perfectly well for this experiment.

Chapter 3

Atom Cloud Transportation

3.1 The Need for Long-Range Transportation

Transportation of laser-cooled atoms over a large distance is a valuable tool for many experiments in atomic physics. Combined with a Bose-Einstein condensate, one can make an ‘atom-laser’, a coherent matter wave that extends over a long distance in space [54, 55, 56]. If an atom is used as a qubit, the ability to move the qubit around for subsequent processing can greatly facilitate the quantum computation, and therefore one could envision the use of the atom transportation as a ‘qubit conveyor belt [57, 30]’.

In my experiment, atom transportation on the atom chip serves to increase the coupling strength (i.e. number of atoms that are coupled) and the duration of the interaction between the thermal atom cloud and the cavity mode. Figure 3.1 shows how the transportation function of the atom chip is utilized for the experiment. The atoms are first collected in a MOT on the atom chip surface, for which a large mirror area is critical for increasing the trapping volume of the MOT and thus increasing the

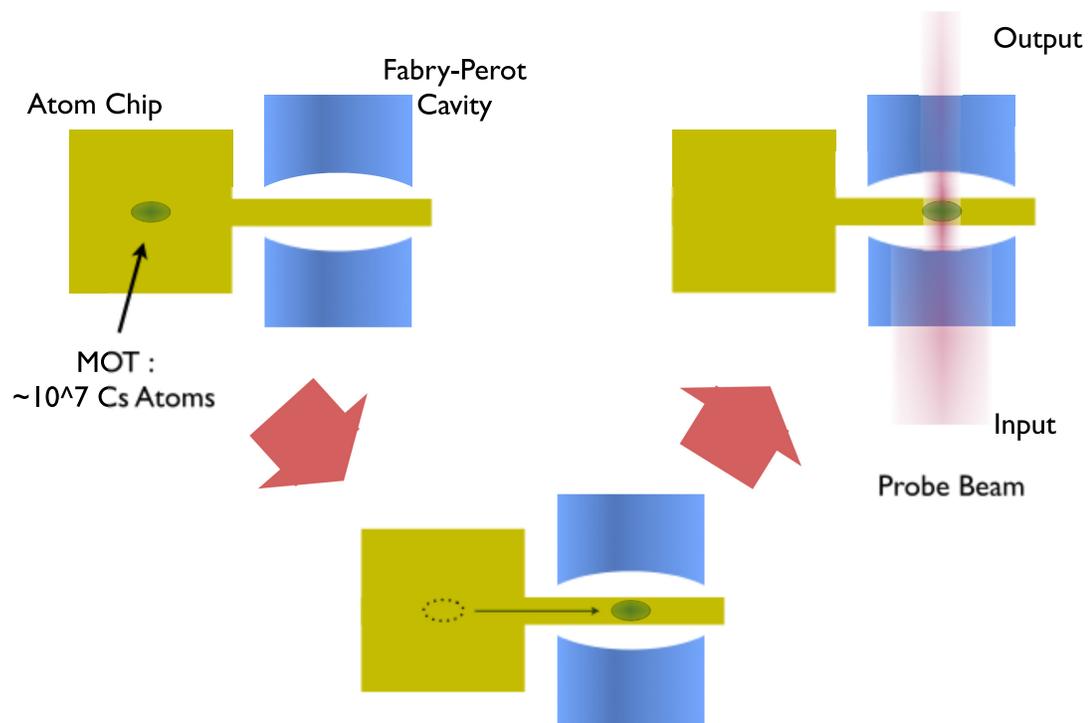


Figure 3.1: A cartoon showing how the atom cloud transportation is used in the experiment.

number of atoms cooled and trapped. These atoms are then transported to the center of a Fabry-Perot cavity magnetically, where they are coupled to the cavity mode and probed by the beam.

Without the use of transportation, coupling a large number of laser-cooled atoms to a cavity mode of such small mode volume is difficult, because you cannot directly create the MOT inside the cavity. Creating the MOT above the cavity and dropping it into the center of the cavity using gravity can also solve this problem, but in that case the atoms pick up speed as they fall through and the transit time for each individual atom is short (< 1 ms) and therefore keeping the atom number constant over the

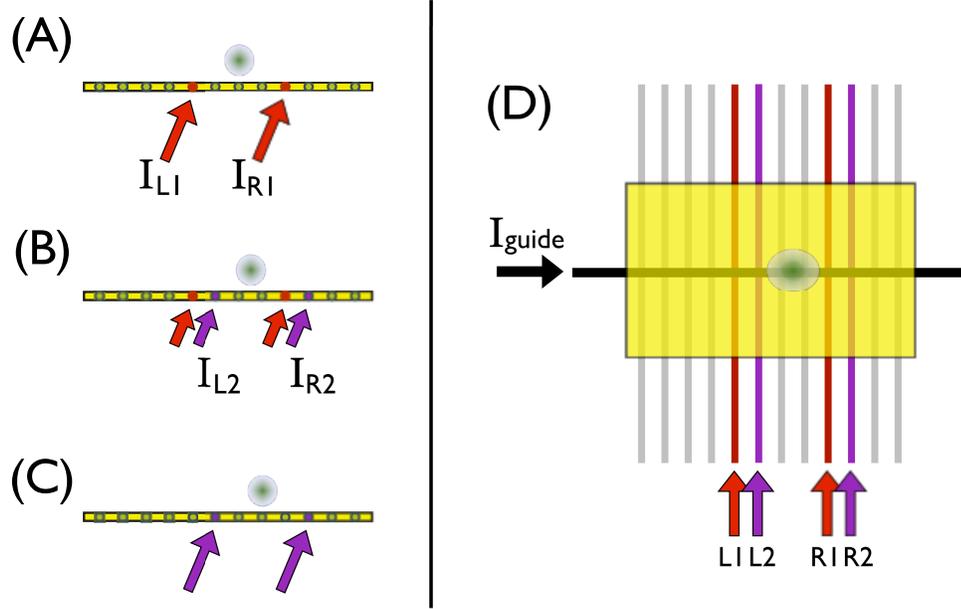


Figure 3.2: Change of currents on transverse wires that transports the atom cloud. (A) to (C) show the sequence of the transportation. (D) shows the top view of the atom chip when the atom cloud is in (B) position.

measurement period becomes a problem [19].

3.2 Transportation Principles

In order to transport atoms along the guide wire, we continuously change the electric currents through the transverse wires to shift the center of the magnetic trap. An illustration of how this is done is shown in Fig. 3.2. The location of the magnetic trap in the guide/transverse wire configuration is at the center of the transverse wire pair. Therefore, in order to move the atom cloud from the location in (A) of Fig. 3.2 to that in (C), we gradually decrease the currents I_{L1} and I_{R1} and at the same time gradually increase I_{L2} and I_{R2} . Such shift of an active transverse wire pair by one

unit distance is repeated until the atom cloud reaches the desired location. In the final atom chip design for my experiment, the unit distance was 1 mm, and the total distance that the atom cloud was transported was 20 mm.

3.3 Some Technical Details

3.3.1 Transportation Speed

In order to transport atoms adiabatically, the speed at which the trap moves during transportation must be slow compared to the the trap frequency in the direction of the transportation. (This direction happens to be the axis where the trapping force is the weakest. The magnetic trap is much tighter in the axes perpendicular to the transportation axis.) In the transportation of Bose-Einstein condensates [31, 30], such adiabatic transportation is critical because it ensures that the heating during the transportation is minimized and the number of atoms in the ground state of the trap is conserved [58].

In our experiment, the trap is very weak (about 10 to 20 Hz in the weak axis) so that the atoms inside the trap are usually far from being in the ground state. When the trap is stationary, we observe the atom cloud oscillating back and forth inside the trap. Fortunately, the temperature of the atoms is low enough that the cloud does not disperse and fill the whole trap during such oscillations. This is very important, because we can use this collective movement of the atom cloud to optimize the transportation.

There are many reasons to prefer faster transportation. First, the number of atoms inside the trap will decay over time due to collisions with themselves (atoms

make transition into an untrappable state) or with the fast-moving background cesium atoms (atoms gain momentum and get kicked out of the trap). Such decay rates are on the order of seconds, so it is best to have the transportation finished before we lose a significant portion of atoms. Second, a slow transportation means that we need to keep the magnetic trap on continuously for a longer time. This can overheat the micro-wires in the atom chip and potentially destroy them. Third, when the experiment is repeated continuously for data taking, faster transportation will shorten the wait time between the measurements.

So, after some trial and error, we abandoned the adiabatic transportation, which would be very slow in our current setup, taking many seconds to even tens of seconds to make the complete delivery. Instead, we chose to actively make use of the collective speed of the atom cloud from a non-ground state. When we make the initial magnetic trap around the atoms (already cooled and optically pumped to the trappable states), we intentionally displace the center of the trap from the center of the atom cloud so that the atoms start rolling down towards the center of the trap. When the atom cloud comes down the trap center, reaching its maximum speed, we start moving the trap at the matched speed. From this point on, the atoms stay near the center of the moving trap until the transportation is finished.

A perfect match between the collective speed of the atom cloud and the speed of the moving trap is not always possible. During most of the transportation attempts, we see that the atom cloud keeps oscillating within the moving trap due to the mismatch of those two speeds. However, as long as the match is close enough, the transportation can be done successfully with minimal atom loss regardless of the presence of the oscillation. Therefore, a fine tuning of the transportation speed is not

required, making the experiment a little less challenging. The transportation speed we use is around 0.1 m/s. The oscillation of the atom cloud due to the mismatch, however, affects the final position of the atom cloud within the magnetic trap at end of the transportation sequence. This needs to be corrected by slightly varying the transportation speed near the end of the transportation sequence.

3.3.2 End of the Transportation Sequence

For the experiments I describe in this thesis, it is important that the atoms are as stationary as possible while they are coupled to the cavity mode. One reason why we want stationary atoms is because we do not take the atom movements into account in our theoretical analysis (Chapter 4). It is possible that the movement of the atoms may cause a change in the measurement, the degree of which has not been extensively studied yet. Another, more practical reason to use nearly-immobile atoms is that this reduces the fluctuation in the number of the coupled atoms during the measurement period, allowing us to make an approximation that the atom number is constant.

To get the most stationary atoms at the end of the transportation sequence, we reverse the sequences we used at the beginning of the transportation. We abruptly stop the movement of the magnetic trap at the end, so that the atom cloud starts climbing up the trap and loses its momentum. The position and the tightness of the final trap should be adjusted so that when the atoms reach the maximum potential energy and lose all their (collective) kinetic energy, they also arrive at the exact center of the cavity mode. The trap z-position (the distance from the atom trap to the chip surface) must also be adjusted to match the height of the cavity mode.

Besides forcing the atoms come to a stop at the end of the transportation, it is

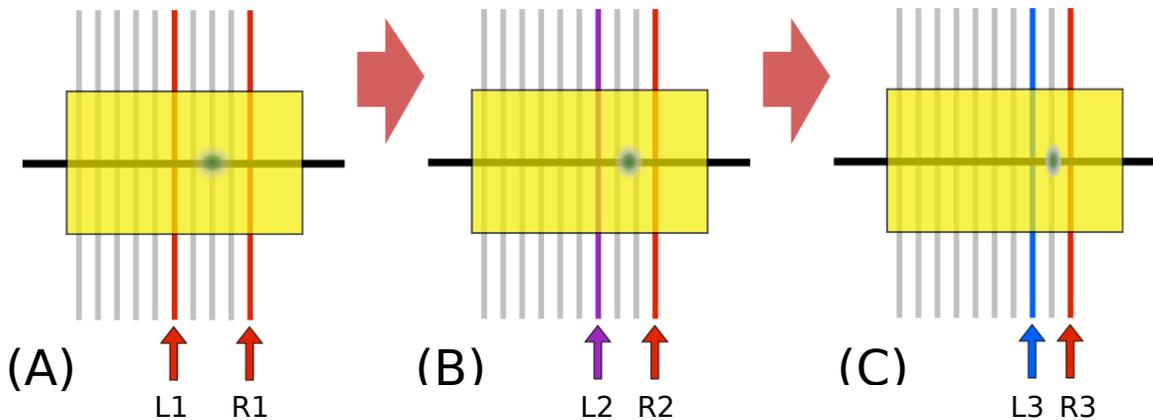


Figure 3.3: Compression sequence at the end of the atom cloud transportation.

also desirable that we maximize the number of atoms coupled to cavity mode¹. One trick we use to boost the number of the coupled atoms is illustrated in Fig. 3.3. Once we get to the final magnetic trap (Fig. 3.3(A), made with a pair of transverse wires 4 mm apart), we fix the current that creates the right wall of the trap (R1, R2 and R3) while we continuously change current for the left wall (L1, L2, and L3), bringing the two walls closer. This effectively compresses the atom cloud and increases the cloud density by a factor of two or more.

3.3.3 Fine-Tunings of the Transportation

Transportation over One Unit Distance

Back to Figure. 3.2, what would be the optimal way to change from full currents in I_{L1} and I_{R1} to full currents in I_{L2} and I_{R2} ? The most straightforward way is to linearly decrease the currents $I_{L1\&R1}$ in time from the full power to zero while at the

¹Once we establish the maximum atom number, we can then easily decrease the number of atoms to a desired level by decreasing the initial MOT size or de-optimizing the transportation sequence.

same time we linearly increase the currents $I_{L2\&R2}$ from zero to the full power. We can simulate the trap movement under such transportation scheme, using the same MATLAB code that generated Fig. 2.8. The simulation result shows that such linear shift in the currents does not lead to a linear translation of the trap center. If the trap movement speed is not constant, the atom cloud inside the trap will experience shakes during their travel.

Using the result of such simulation, it is actually possible to reverse-engineer how the weight should be shifted between $I_{L1\&R1}$ and $I_{L2\&R2}$ as a function of time that will ensure the translation of the trap center at a constant velocity. This new function can be improved even more, because the trap frequency can also change during the transition from $I_{L1\&R1}$ to $I_{L2\&R2}$. An overall factor that applies equally to both $I_{L1\&R1}$ and $I_{L2\&R2}$ can be designed to correct this problem, too.

I worked out such an optimized function for the current shifting operation and applied it to the experiment. The improvement in efficiency of atom transportation was not found, however. It seems that the transportation already works pretty well regardless of the imperfection of the shifting operation. So, in the end, I simply used the linear shift in the currents, as it was the easiest to implement.

Transportation over the Total Distance

As is described in [58]², there could be better ways to fast transport atoms than using a constant speed. Why not gradually increase the transportation speed for the first half and gradually decrease the speed for the last half? I tried many variations of such transportation trajectories such as a constant acceleration/deceleration scheme

²The referenced paper cannot be directly applied to our experiment, because the magnetic trap we use is much too weak and therefore achieving the vibrational ground state is never our goal.

or a Gaussian error function. In the end, however, I could not find any significant improvements and I stuck with the fixed-speed transportation.

3.3.4 Scattered Light

While the atoms are being transported by the magnetic trap, we need to make sure that there is absolutely no stray light that can get into the vacuum chamber and excite the atoms in the trap. During the early days when we were testing the transportation, the “scattered light” from regular optical components was found to significantly reduce the lifetime of the atom cloud. This discovery was made somewhat by luck while I was looking for the cause of fast decay of atom number in the magnetic trap. I had never suspected that such low-powered scattered (not directed) light could have any effect on the experiment. Now we have barriers between the major optical components on the optical table and the vacuum chamber that will block any scattered light.

During the transportation of the atom cloud, the MOT beams are no longer needed and have to be turned off. Before we realized how sensitive the atoms clouds in the magnetic trap can be to the ambient light, the tapered amplifier that creates the MOT beams used to be controlled solely by switching its seeding beam. Now we have an extra beam shutter (made with an ordinary speaker [49]) that blocks the weak beam that comes out of the tapered amplifier even at no seeding. All these protections have been verified to increase the lifetime of the atom cloud significantly.

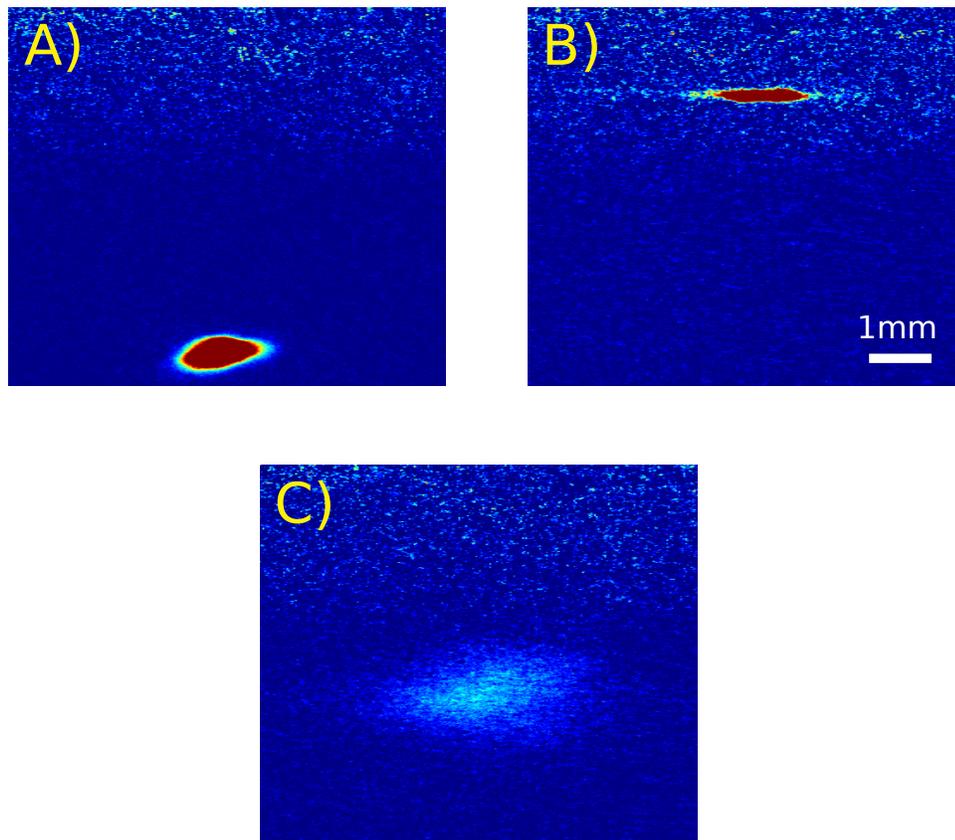


Figure 3.4: Fluorescence images of an atom cloud. (A) Macro-MOT. (B) Micro-MOT. (C) PG-cooled atom cloud after 20 ms free fall from the micro-MOT position.

3.4 Atom Cloud Imaging

3.4.1 Fluorescence Imaging

In our experiment, fluorescence imaging is used as a main tool to examine the status of the atom cloud before the transportation begins.

Macro-/Micro-MOT Imaging for Atom Number Estimation

In the current experiment setup, the MOT is constantly monitored by a real-time video camera so that we can use it to check the MOT shape and optimize its size. But we also have a separate scientific-grade CCD camera with a mechanical shutter (Model : Photometrics SenSys) that digitally records the MOT image for detailed analysis. This camera is positioned right next to one of the side windows of the vacuum chamber and looks at the cloud underneath the atom chip from the side.

Imaging a MOT is very easy, because the atoms in the MOT are constantly excited by the MOT beam and naturally fluoresce; it requires no additional illumination for imaging. Assuming that the atoms spend half the time in the excited state, we can convert the CCD count integrated over the exposure period into the number of photons it received and estimate the atom number inside the MOT. An example of a MOT image is shown in Figure 3.4 (A).

Initially, the MOT is formed at the center of the cooling/trapping region, quite a distance from the atom chip surface. By reducing the current that creates the required magnetic field gradient for the MOT, the position of the MOT can be temporarily lowered,³ closer to the surface, so that the MOT position spatially overlaps with the magnetic trap. We call this near-surface MOT the ‘micro-MOT’. The fluorescence image of the micro-MOT is used to confirm the good transfer of the atoms from the initial MOT (or sometimes called ‘macro-MOT’ to make the distinction between the two more clear). An example of the micro-MOT is shown in Figure 3.4 (B).

³Or ‘raised’ to be exact, since the atom chip is installed inside the vacuum chamber upside down.

Imaging After PG Cooling to Check the Temperature

Once the atoms are brought close to the surface by the micro-MOT, PG cooling [59] is applied to lower the atom temperature down to $\sim 10\mu\text{K}$. To confirm that the cooling was successfully accomplished, pictures of the atom cloud in free-fall are taken and the spread of the cloud is analyzed. By measuring the speed of the spread, it is possible to deduce the temperature of the cloud. In practice, when we quickly want to confirm that the PG cooling has worked, we take a photo after ~ 20 ms from the release and check if the atom cloud with a similar size of the micro-MOT is visible. Without the proper cooling, the atom cloud at its initial Doppler-limited temperature $\sim 100\mu\text{K}$ would spread so fast that after 20 ms a very thin cloud many times larger than the original micro-MOT size would be (barely) visible when imaged. An example of a PG-cooled atom cloud after 20 ms freefall is shown in Figure 3.4 (C).

To record a fluorescence image of the atom cloud, unlike the imaging of the MOT, we must first provide illumination so that the atoms can fluoresce. We re-use the MOT beams for this purpose, which is very convenient because then we do not need to prepare a new separate light source for the imaging. The MOT beams have, by default, a good power balance between the pairs of counter-propagating beams, which helps to minimize the movement of the atoms during the exposure. We use from 5 to 20 milliseconds for the exposure.

One thing to watch out for when imaging atoms in free-fall is that if there is a non-zero magnetic field at the cloud location, the atoms may experience a force from the excitation beam and be pushed away during the exposure period. This makes the atom image blurred, or in the worst case there will barely be any atom image to see. To prevent this, the magnetic field must be zeroed before imaging. Fortunately, PG

cooling already requires zero magnetic field to work, so once we figure out what bias field we need to apply to cancel the ambient field and have the PG cooling working, we can apply the same magnetic field settings for the fluorescence imaging as well.

Imaging of trapped atoms

After PG cooling, we apply optical pumping to the atoms to drive them into the weak-field-seeking state, and then we create the quadrupole magnetic trap around them using currents through the micro-wires in the atom chip and the bias field from the Helmholtz coils surrounding the vacuum chamber [45, 46]. To test the trap performance, we take pictures of the atom cloud in the trap after certain time periods and check how fast the atom number in the trap decays. Normally when everything is in order, we would see the cloud moving back and forth along the weak trapping axis with frequency of 50-200 Hz and the number of atoms would decay with a half-life of 1 second or more⁴.

Trapped atom imaging is primarily used to check whether the magnetic trap works properly, but it is also used as importantly to see where the initial position of the atom cloud is relative to the trap center. If these two coincide when the trap is formed, the oscillation of the atom cloud within the trap would not be visible. We actually want the atom cloud to start away from the center of the trap, as explained in subsection 3.3.1, so that it can gain momentum as it moves towards the cavity. By imaging the atom cloud movement during the first few tens of milliseconds and watching which direction it moves and how it oscillates within the magnetic trap, we can determine where the magnetic trap center is. Using this knowledge, we can

⁴I almost never try the magnetic trapping for more than a few hundred of milliseconds, just to be safe.

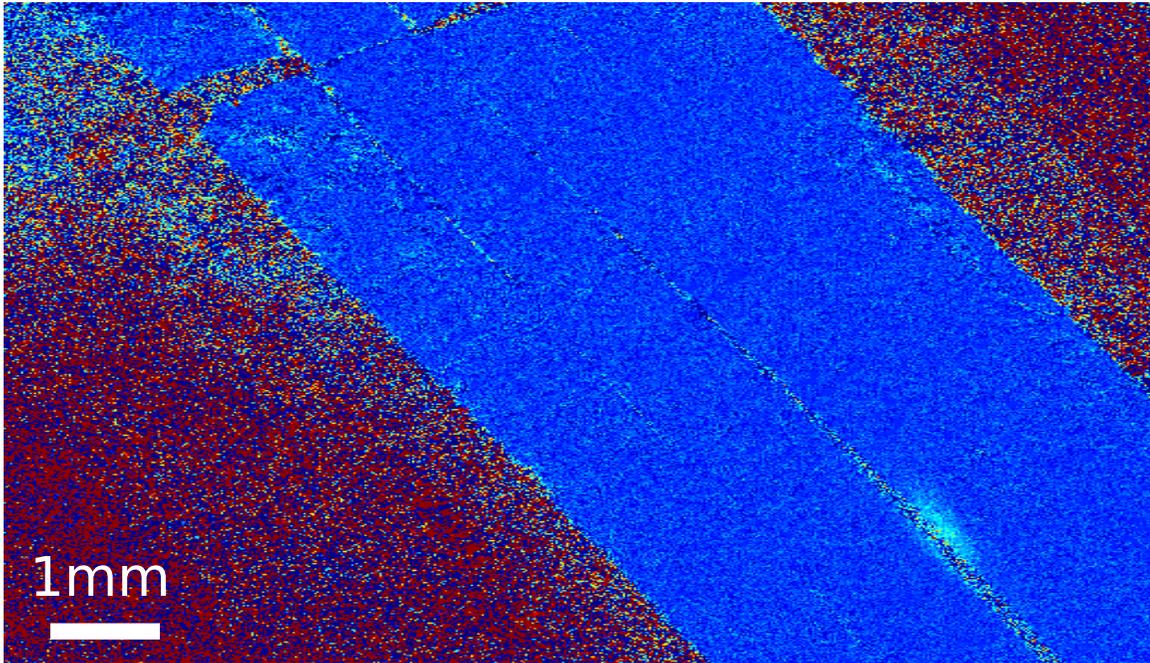


Figure 3.5: Absorption image of the atom cloud after it was transported to the center of the cavity. The bright spot on the lower right is the atom cloud, and the blue area is the mirror surface of the narrow channel of the atom chip that extends into the cavity center (Fig. 3.1). Parallel lines visible on the atom chip surface are unwanted (but harmless) cracks on the mirror that occurred after the guiding wires were inserted on the backside of the atom chip. When imaged at different times during the transportation, we see the atom cloud moving along the central crack line from the top left to the bottom right.

adjust the micro-MOT position and change the starting position of the atom cloud accordingly.

3.4.2 Absorption Imaging

Once the transportation starts, we use absorption imaging to follow the movement of the atoms along the narrow channel between the two cavity mirrors. Fluorescence imaging cannot be used in this case because the optical access for the MOT beams

is blocked by the Fabry-Perot cavity mirrors. Without the red-detuned (Doppler cooling) MOT beams from all six directions, the atom cloud cannot maintain its position for the tens of milliseconds that we need for the fluorescence imaging, and it will fall down due to gravity or being pushed away by the radiation pressure during the exposure. This forces us to use the absorption imaging instead, which is more sensitive than the fluorescence imaging. We need less than a millisecond exposure to get a decent photo of the atom cloud, an example of which is shown in Fig. 3.5.

Unlike the fluorescence imaging camera that is positioned on the side of the vacuum chamber (so that the atom chip is imaged as a horizontal line), the absorption imaging camera (Model Name: pixelfly from PCO) views the atom chip from the bottom. This is possible thanks to the mirror surface of the atom chip. The absorption imaging beam, after it passes through the atom cloud, gets reflected on the surface of the chip and is directed into the pixelfly camera.

3.5 Controlling the Currents

As explained in Section 3.2 and Fig. 3.2, the transportation of the atom cloud is made possible by the use of many underlying transverse micro-wires. In our latest atom chip design, more than two dozen micro-wires were used to move the atoms over the distance of 2 cm. Independently controlling currents in so many wires poses some simple but rather technically challenging problems.

The first challenge comes from the fact that it would be impractical to use a separate power supply for each micro-wire. Preparing more than two dozen power supplies would cost a lot, they would take up too much lab space, and this would require too many ADC ports from the control computer. The second problem is, even

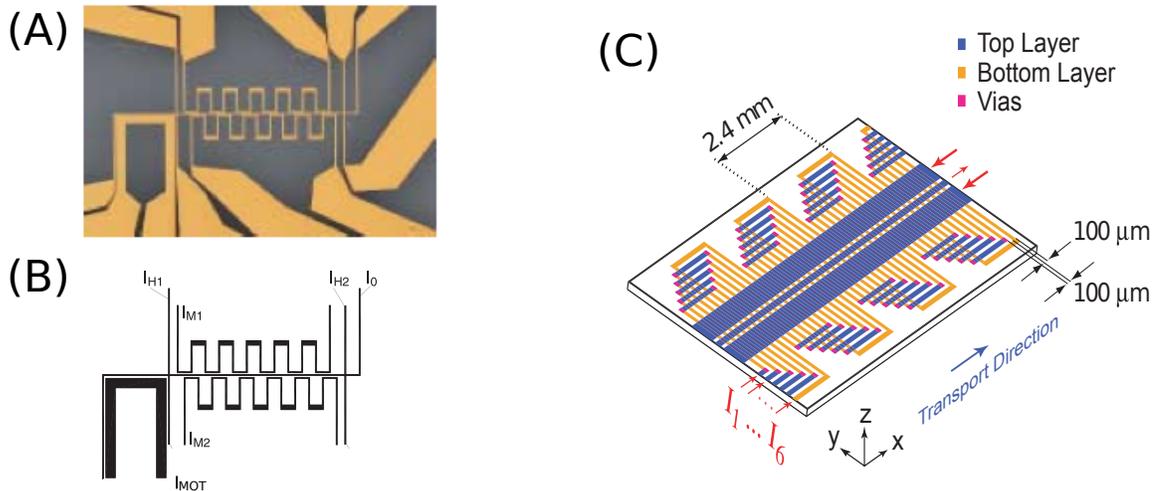


Figure 3.6: Atom chip designs from Reichel Group in Kastler Brossel Laboratory, France. (A) A photograph of the atom chip substrate (before mirror coating) with transverse wires in zig-zag formation, copied from [60]. (B) Schematics of the wire layout of (A). (C) Two-layer atom chip design with 6 separate transverse wire patterns, that was used to transport atoms over 10 cm [61].

if we had the power supply problem under control, there is a practical limit on the number of connector pins we can place on the PCB to connect to the power supplies. In addition to the problem of limited space on the PCB, it is also difficult to fit so many wires in feedthroughs to the vacuum chamber.”

Clever atom chip designs that are worth noting that circumvent such problems can be found from the experiments of the Reichel Group [60, 61]. Some examples are shown in Fig. 3.6. The key idea in their design is that they connect transverse wires together so the one power supply can be used run currents through multiple transverse wires simultaneously. This works because the atom cloud care only about the field gradients coming from the nearest wires.

Our experiment uses a more direct approach and uses a relay box to choose the

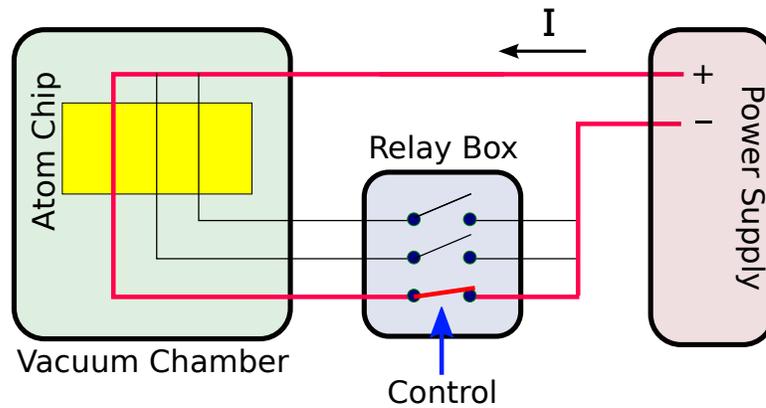


Figure 3.7: A schematic diagram explaining how the relay box is used to route the current from a power supply to the chosen transverse wire. In the actual experiment, four independent relay boxes are used, each controlling up to 16 switches.

transverse wire to run currents through from a single power supply. (See Figure 3.7.) Four independent power supplies are used to drive four transverse wires (I_{L1} , I_{R1} , I_{L2} , and I_{R2} in Fig. 3.2) that are necessary to shift the magnetic trap by one unit distance (1 mm). After a unit-distance transportation is finished, control signals (as in Fig. 3.7) into the relay boxes are changed, which re-defines I_{L1} through I_{R2} , and a new cycle of a unit-distance transportation can begin.

This design allowed us to cut down the needed number of power supplies to only four⁵ for the atom cloud transportation to an arbitrarily long distance. Also, all the transverse wires that share the same power supply were combined on one end (the side where it is not connected to the relay box) directly on the PCB board level, saving us the number of the connector pins and the feedthrough connections by a half. For both PCB and the vacuum chamber feedthroughs, two 25-pin D-Sub connectors were

⁵In order to run the compression sequence at the end of the transportation (subsection 3.3.2 and Fig. 3.3), the last transverse wire is controlled by a separate power supply, making the total number of power supplies used to control transverse wires to five.

used to make all the necessary electrical connections.

Chapter 4

Semiclassical Approach to Multi-Atom Cavity QED Theory

In this chapter I explain the semiclassical theory that I use to analyze my experimental results. I also use this theory to run simulations and search for the better parameters for the experiment. The key equations of motion in this theory is what we casually¹ call the Maxwell-Bloch equations (MBE's). We first start with a simplest model, a Fabry-Perot cavity with a single two-level atom inside, and derive the MBE's (Section 4.1). We then generalize the equations to take account of multiple atoms inside the cavity, a case that is experimentally more relevant (Section 4.2). In the subsequent sections, we explain how the theory is applied to the experiment.

¹I say 'casually', because it is not the most appropriate use of the name. The original use of the term [62] simply refers to the combination of two separate equations: one is the the Maxwell's (macroscopic) equations for which the atomic vapor is considered as a dielectric material. They describe the evolution of the electromagnetic field in the classical framework. The other is the optical Bloch equation, that describes the evolution of the atom's quantum state due to the external field. For a more detailed analysis of the Maxwell-Bloch equations, refer to [2].

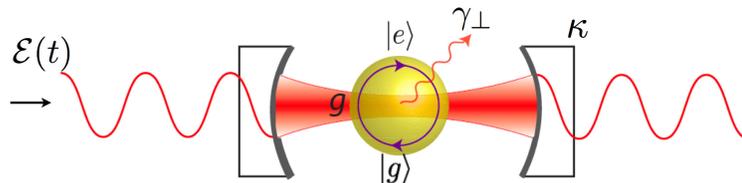


Figure 4.1: A cartoon of a single two-level atom and a Fabry-Perot cavity system.

4.1 Maxwell-Bloch Equations for a Single Atom in a Cavity

The model we use in this first section is depicted in Figure 4.1. A two-level atom is placed inside a Fabry-Perot cavity and is driven by an injection laser. A photon that is injected into the cavity can be absorbed by the atom in the ground state, bringing it up to the excited-state. Similarly, the excited-state atom can re-emit the photon into the cavity mode and fall down to the ground-state. Such exchange of photons between the cavity mode and the atom is characterized by the coupling strength g .

We assume that there are two decay channels of the system, as also depicted in the figure. The first channel is the spontaneous emission, by which an excited atom emits a photon into the free-space, rather than into the cavity mode. The rate at which the system loses a photon by this mechanism is $2\gamma_{\perp}$. The second type of decay happens at the cavity mirrors that are assumed to be not 100% reflective. Some photons will leak through the mirrors and escape from the system through the cavity output mode at the rate of 2κ . The energy inside the system can be replenished by an external drive (i.e. input laser), the strength of which is denoted by \mathcal{E} .

The Jaynes-Cummings Hamiltonian for the atom-cavity system with the driving

term is given by

$$\mathbf{H} = \Delta_C \mathbf{a}^\dagger \mathbf{a} + \Delta_A \boldsymbol{\sigma}_+ \boldsymbol{\sigma}_- + ig (\mathbf{a}^\dagger \boldsymbol{\sigma}_- - \mathbf{a} \boldsymbol{\sigma}_+) + i\mathcal{E} (\mathbf{a}^\dagger - \mathbf{a}) \quad (4.1)$$

where Δ_C is the cavity detuning ($\omega_{\text{cavity}} - \omega_{\text{laser}}$) and Δ_A is the atom detuning ($\omega_{\text{atom}} - \omega_{\text{laser}}$). The symbols for the operators follow the usual conventions: \mathbf{a} and \mathbf{a}^\dagger are the intracavity field annihilation and creation operators, and $\boldsymbol{\sigma}_-$ and $\boldsymbol{\sigma}_+$ are atom lowering and raising operators.

The evolution of the system, expressed by a (reduced) density operator $\boldsymbol{\rho}$, is then given by the master equation

$$\frac{d}{dt} \boldsymbol{\rho} = -i [\mathbf{H}, \boldsymbol{\rho}] + \gamma_\perp (2\boldsymbol{\sigma} \boldsymbol{\rho} \boldsymbol{\sigma}^\dagger - \boldsymbol{\sigma}^\dagger \boldsymbol{\sigma} \boldsymbol{\rho} - \boldsymbol{\rho} \boldsymbol{\sigma}^\dagger \boldsymbol{\sigma}) + \kappa (2\mathbf{a} \boldsymbol{\rho} \mathbf{a}^\dagger - \mathbf{a}^\dagger \mathbf{a} \boldsymbol{\rho} - \boldsymbol{\rho} \mathbf{a}^\dagger \mathbf{a}). \quad (4.2)$$

Now, we can derive equations of motion for expectation values of any operators, using the rule

$$\frac{d}{dt} \langle \mathbf{O} \rangle = \frac{d}{dt} \text{Tr} [\boldsymbol{\rho} \mathbf{O}] = \text{Tr} \left[\frac{d\boldsymbol{\rho}}{dt} \mathbf{O} \right] \quad (4.3)$$

where \mathbf{O} is an arbitrary operator. Applying Eq. 4.3 to operators \mathbf{a} , $\boldsymbol{\sigma}_-$ and $\boldsymbol{\sigma}_z$, we can easily derive [27]

$$\langle \dot{\mathbf{a}} \rangle = -(\kappa + i\Delta_C) \langle \mathbf{a} \rangle + g \langle \boldsymbol{\sigma}_- \rangle + \mathcal{E}(t) \quad (4.4a)$$

$$\langle \dot{\boldsymbol{\sigma}}_- \rangle = -(\gamma_\perp + i\Delta_A) \langle \boldsymbol{\sigma}_- \rangle + g \langle \mathbf{a} \rangle \langle \boldsymbol{\sigma}_z \rangle \quad (4.4b)$$

$$\langle \dot{\boldsymbol{\sigma}}_z \rangle = -2\gamma_\perp (\langle \boldsymbol{\sigma}_z \rangle + 1) - 2g (\langle \boldsymbol{\sigma}_+ \rangle \langle \mathbf{a} \rangle + \langle \boldsymbol{\sigma}_- \rangle \langle \mathbf{a}^\dagger \rangle) \quad (4.4c)$$

where $\boldsymbol{\sigma}_z \equiv |e\rangle\langle e| - |g\rangle\langle g| = [\boldsymbol{\sigma}_+, \boldsymbol{\sigma}_-]$. This set of equations are what we in the Mabuchi Group normally call the Maxwell-Bloch equations.

The ‘factorization approximation’ was used during the derivation for Eqs. 4.4, where the expectation values of products of two operators are substituted by products of two expectation values. For example

$$\langle \mathbf{a}\boldsymbol{\sigma}_z \rangle \rightarrow \langle \mathbf{a} \rangle \langle \boldsymbol{\sigma}_z \rangle. \quad (4.5)$$

This approximation means that any entanglement between the atom and the field is ignored. There are ways to add higher order terms to improve the accuracy of the approximation, but in general this first order approximation is enough to capture most of the interesting behaviors we expect to see from the system.

Equations 4.4 can also be written in a slightly different, simpler form, by introducing the new variables

$$\Theta \equiv \frac{\Delta_C}{\kappa}, \quad \Delta \equiv \frac{\Delta_A}{\gamma_\perp} \quad (4.6)$$

and

$$x \equiv \frac{\sqrt{2}g}{\gamma_\perp} \langle \mathbf{a} \rangle, \quad p \equiv -\sqrt{2} \langle \boldsymbol{\sigma}_- \rangle, \quad d \equiv -\langle \boldsymbol{\sigma}_z \rangle, \quad y \equiv \frac{\sqrt{2}g}{\kappa\gamma_\perp} \mathcal{E}. \quad (4.7)$$

Equations 4.4 are then now written as²

$$\dot{x} = -\kappa \left\{ (1 + i\Theta)x + 2Cp - y \right\} \quad (4.8a)$$

$$\dot{p} = -\gamma_\perp \left\{ (1 + i\Delta)p - xd \right\} \quad (4.8b)$$

$$\dot{d} = -2\gamma_\perp \left\{ d - 1 + \frac{x^*p + xp^*}{2} \right\} \quad (4.8c)$$

²The reason I factor out $2\gamma_\perp$ in Eq. 4.8c rather than γ_\perp is because it can be replaced by γ_\parallel later when we need to modify the equations under the condition $\gamma_\parallel \neq 2\gamma_\perp$ (section 4.4).

where C is the cooperativity parameter, defined as

$$C \equiv \frac{g^2}{2\kappa\gamma_\perp}. \quad (4.9)$$

4.2 Generalization to Multi-Atom Case

Now, we assume that there are multiple atoms inside the cavity, all coupling with different coupling strengths. To be more specific, we assume that there are N atoms, and the j -th atom has a coupling strength g_j such that

$$g_j = c_j \cdot g_0 \quad (j = 1, 2, 3, \dots, N) \quad (4.10)$$

where g_0 is the maximum coupling strength occurring at the center of the cavity, and c_j is the coupling coefficient that takes values between 0 and 1.

The generalized Jaynes-Cummings Hamiltonian for multi-atom case is

$$\mathbf{H} = \Delta_C \mathbf{a}^\dagger \mathbf{a} + \sum_{j=1}^N g_j \Delta_A \boldsymbol{\sigma}_+^j \boldsymbol{\sigma}_-^j + i \sum_{j=1}^N g_j \left(\mathbf{a}^\dagger \boldsymbol{\sigma}_-^j - \mathbf{a} \boldsymbol{\sigma}_+^j \right) + i\mathcal{E} (\mathbf{a}^\dagger - \mathbf{a}) \quad (4.11)$$

and the corresponding Maxwell-Bloch equations are

$$\langle \dot{\mathbf{a}} \rangle = -(\kappa + i\Delta_C) \langle \mathbf{a} \rangle + \sum_{j=1}^N g_j \langle \boldsymbol{\sigma}_-^j \rangle + \mathcal{E}(t) \quad (4.12a)$$

$$\langle \dot{\boldsymbol{\sigma}}_-^j \rangle = -(\gamma_\perp + i\Delta_A) \langle \boldsymbol{\sigma}_-^j \rangle + g_j \langle \mathbf{a} \rangle \langle \boldsymbol{\sigma}_z^j \rangle \quad (4.12b)$$

$$\langle \dot{\boldsymbol{\sigma}}_z^j \rangle = -2\gamma_\perp (\langle \boldsymbol{\sigma}_z^j \rangle + 1) - 2g_j (\langle \boldsymbol{\sigma}_+^j \rangle \langle \mathbf{a} \rangle + \langle \boldsymbol{\sigma}_-^j \rangle \langle \mathbf{a}^\dagger \rangle) \quad (4.12c)$$

where the superscript j on sigma operators means the operators for the j -th atom.

Using the definition

$$x \equiv \frac{\sqrt{2}g_0}{\gamma_\perp} \langle \mathbf{a} \rangle, \quad p_j \equiv -\sqrt{2} \langle \boldsymbol{\sigma}_-^j \rangle, \quad d_j \equiv -\langle \boldsymbol{\sigma}_z^j \rangle, \quad y \equiv \frac{\sqrt{2}g_0}{\kappa\gamma_\perp} \mathcal{E}, \quad (4.13)$$

Eqs. 4.8 can be generalized to

$$\dot{x} = -\kappa \left\{ (1 + i\Theta)x + 2 \sum_{j=1}^N c_j C_0 p_j - y \right\} \quad (4.14a)$$

$$\dot{p}_j = -\gamma_\perp \left\{ (1 + i\Delta)p_j - c_j x d_j \right\} \quad (4.14b)$$

$$\dot{d}_j = -2\gamma_\perp \left\{ d_j - 1 + c_j \frac{x^* p_j + x p_j^*}{2} \right\} \quad (4.14c)$$

with the cooperativity C_0 now defined based on the maximum coupling strength (g_0) for a single atom,

$$C_0 \equiv \frac{g_0^2}{2\kappa\gamma_\perp}. \quad (4.15)$$

Note that in Eq. 4.17a, the term $c_j C_0 p_j$ may seem incorrect, as one might think that it should have been $c_j^2 C_0 p_j$ since C_0 has g_0^2 in it. However, this is not quite so, because as $\langle \mathbf{a} \rangle$ is replaced by x , it absorbs one g_0 into its definition.

Once the values of c_j 's ($j = 1, 2, 3, \dots, N$) are specified, Equations 4.14 will provide us with the equations of motions for the atomic states of individual atoms (from p_j and d_j), and the intracavity field strength (from x). The field amplitude of the cavity output, which is the only quantity we directly measure in the experiment, can also be calculated because it is proportional to the intracavity field strength by a single factor κ .

There is one minor problem with this direct approach, however. The number of atoms (N) that we couple to the cavity in the experiment is on the orders of millions.

The size of the Eqs. 4.14 grows linearly with N and this quickly gets out of hand. In order to make the computation more manageable, we can group the atoms that have similar coupling strengths (c_j 's) and treat them all as behaving identically. This is a good approximation as long as the values of the p_j 's used in Eq. 4.17a do not vary significantly among those in the same group. Note that this condition is easily met, not only when the system has reached a steady state, but also when the system is oscillating in a limit cycle. Even though our model assumes no direct coupling between two atoms, all the atoms are coupled to the same cavity field, and therefore the states of the atoms in motion are automatically synchronized with one another after they interact with the field long enough.

In order to carry out such approximation, we first need to discretize c_j space. We define a set

$$S = \{s_1, s_2, s_3, \dots, s_L\} \quad (4.16)$$

where values of all elements s_k in the set S are between 0 and 1. Then we can define n_k ($k = 1, 2, \dots, L$) as the number of c_j 's that are most close to s_k among all elements in S , and approximate the Eqs. 4.14 by

$$\dot{x} = -\kappa \left\{ (1 + i\Theta)x + 2 \sum_{k=1}^L n_k s_k C_0 p_j - y \right\} \quad (4.17a)$$

$$\dot{p}_k = -\gamma_{\perp} \left\{ (1 + i\Delta)p_k - s_k x d_k \right\} \quad (4.17b)$$

$$\dot{d}_k = -2\gamma_{\perp} \left\{ d_k - 1 + s_k \frac{x^* p_k + x p_k^*}{2} \right\} \quad (4.17c)$$

We have empirically found that with an evenly spaced elements s_k 's, the length L of the set S only needs to be more than 30 for this discretized-coupling-approximation to be fairly accurate. Most of my calculations were done using $L = 300$ just to be safe,

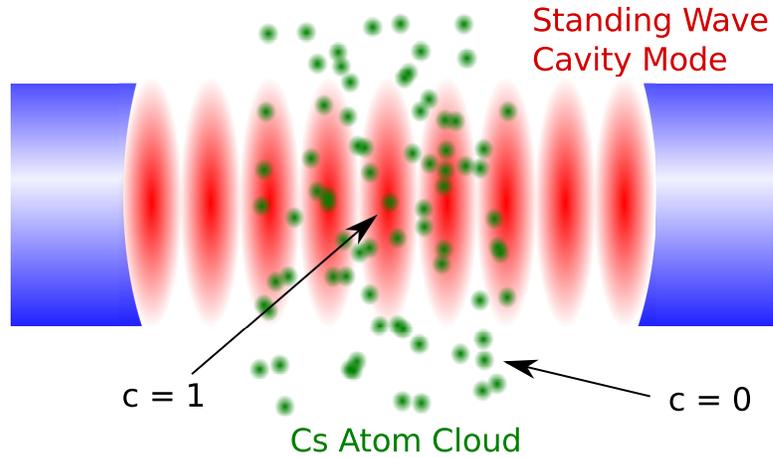


Figure 4.2: A cartoon of an atom cloud coupled to a standing wave cavity mode. An atom that has the largest coupling strength, g_0 , to the cavity field is indicated ($c = 1$). An atom with a coupling strength close to zero is also indicated ($c = 0$).

since even at this level the computation can be handled reasonably well by modern laptop computers.

The determination of the n_k 's, or in other words, the choice of the set $\{c_j\}$ depends on two factors. The first is the atom number N , and the second is how the atoms are distributed inside the cavity mode. N may change from experiment to experiment, but there is a way to estimate the value of N for each experimental run as will be described later in this chapter, Section 4.3.3. The distribution of atoms, on the other hand, cannot be easily measured experimentally. Instead, we simply assume that the atoms are distributed uniformly inside the cavity. This is a reasonable assumption, because the size of the atom cloud (dimensions $> 100\mu\text{m}$) that we transport into the cavity using the atom chip is larger than the size of the cavity mode (waist size $\sim 65\mu\text{m}$).

Given the assumption of the uniform distribution of the atoms in space, I used a

Monte Carlo method to figure out what the expected distribution of n_k is. I randomly placed atoms with equal probability everywhere, and computed the histogram for the values of $c_j (= g_j/g_0)$, where the value of c_j is calculated as a function of the atom position relative to the gaussian standing-wave profile of the TEM₀₀ cavity field mode (See Figure 4.2 and 4.3). Once the distribution of the n_k is found, we can adjust the overall scale of the n_k 's to match the condition $\sum_{k=1}^L n_k = N$, and this completes our search for the values for n_k 's to be used in Equations 4.17.

Figure 4.3 shows the Monte-Carlo simulation result with $\sim 10^6$ samples, I used these results as a base for the all the subsequent theoretical calculations for my experiment. I should also mention that the distribution of the n_k 's can be calculated analytically without the need for the Monte Carlo method³.

4.3 Steady State Solutions

4.3.1 How to Solve for the Steady State

The Equations 4.17 can be used to find the steady state solutions, by setting all the derivatives on the left-hand-side of the equations to zero :

$$0 = (1 + i\Theta)x + 2 \sum_{k=1}^L n_k s_k C_0 p_k - y \quad (4.18a)$$

$$0 = (1 + i\Delta)p_k - s_k x d_k \quad (4.18b)$$

$$0 = d_k - 1 + s_k \frac{x^* p_k + x p_k^*}{2}. \quad (4.18c)$$

³Nikolas in our group worked it out analytically after I showed him my method.

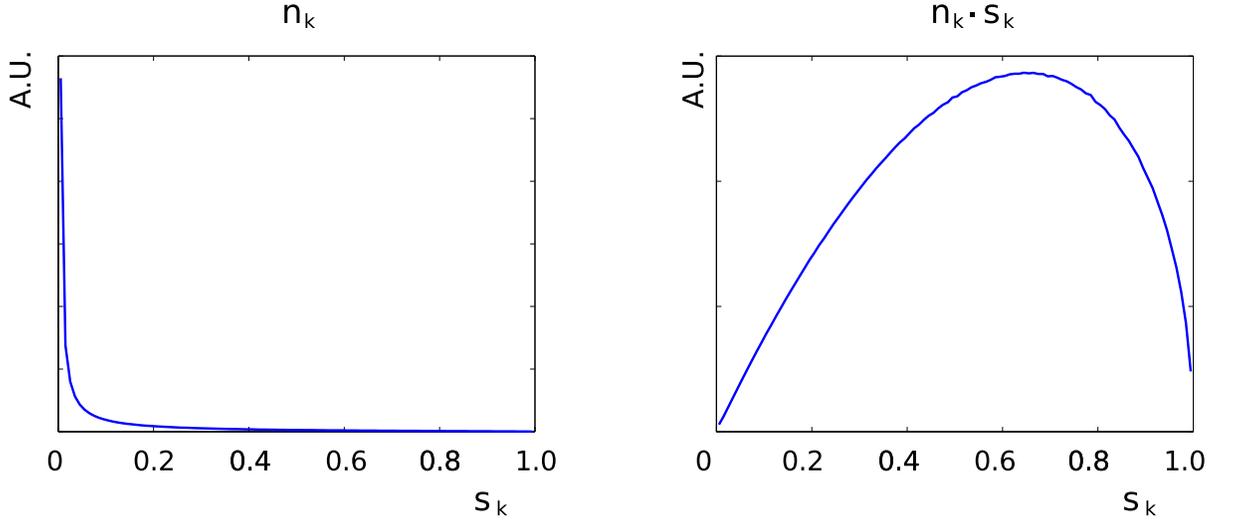


Figure 4.3: Monte Carlo simulation results that I used to model my experiment. On the left, n_k values are calculated for each s_k . Because n_k value diverges at $s_k = 0$, I also calculated $n_k s_k$ as a sanity check, and the result is displayed on the right. The $n_k s_k$ plot is informative, because it shows atoms with what range of s_k play the most crucial rule in our theory.

Now, we want to solve these multi-variable equations and write x , p_k , and d_k as functions of y , the input drive strength. If done in the right order, solving these equations is very easy. First, we can re-write Eq. 4.18b as

$$p_k = \frac{s_k x d_k}{1 + i\Delta} \quad (4.19)$$

and plug it into Eq. 4.18c. This allows us to eliminate the term p_k and get

$$d_k = \frac{1 + \Delta^2}{1 + \Delta^2 + s_k^2 |x|^2}. \quad (4.20)$$

We can plug this d_k back into Eq. 4.19, and get

$$p_k = \frac{s_k x (1 - i\Delta)}{1 + \Delta^2 + s_k^2 |x|^2}. \quad (4.21)$$

And finally, this p_k is plugged into Eq. 4.18a and we obtain

$$y = x \left\{ (1 + i\Theta) + \sum_{k=1}^L \frac{2n_k s_k^2 C_0 (1 - i\Delta)}{1 + \Delta^2 + s_k^2 |x|^2} \right\} \quad (4.22)$$

or since we are mostly interested in the magnitudes,

$$|y| = |x| \sqrt{\left(1 + \sum_{k=1}^L \frac{2n_k s_k^2 C_0}{1 + \Delta^2 + n_k s_k^2 |x|^2}\right)^2 + \left(\Theta - \sum_{k=1}^L \frac{2n_k s_k^2 C_0 \Delta}{1 + \Delta^2 + n_k s_k^2 |x|^2}\right)^2}. \quad (4.23)$$

Equations 4.22, 4.21, and 4.20 constitute the steady state solution of $\{x, p_j, d_j\}$ of our system given y .

One possible source of confusion in understanding these steady state solutions is that they are all written as a function of x (\propto output of the cavity), not as a function of y (\propto input of the cavity). The reason is obvious in the form of Eq. 4.23. Depending on the system parameters $\{\Delta, \Theta, C_0, N\}$, the steady state output x can be a multi-valued function of the input y . One example of such a case is shown in Figure 4.4.

4.3.2 Stability of the Steady State Solutions

Given a steady state solution, as calculated above, we can also determine the stability of this solution. The Jacobian matrix can be used for this purpose. In order to use the Jacobian matrix, one must also include equations of motions for \dot{x}^* and \dot{p}_k^* , because

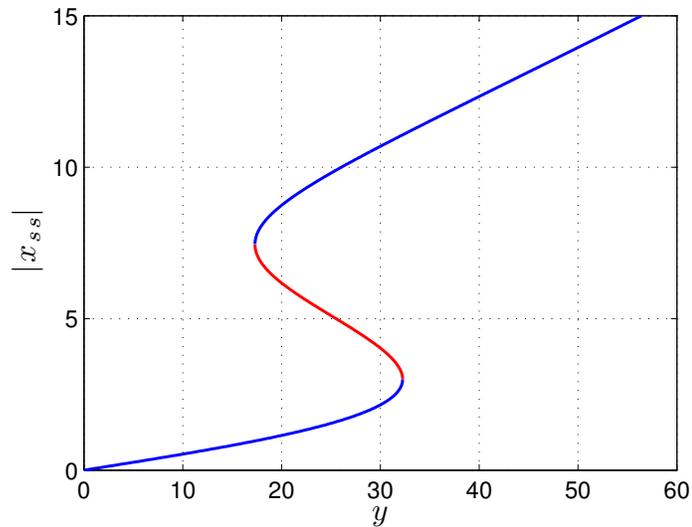


Figure 4.4: A steady state solution of x as a function of y . This is a single atom case ($N = 1$, and therefore $L = n_1 = s_1 = 1$), with $C = 47$, $\Theta = 5.0$, and $\Delta = 3.8$. Blue indicates stable solution, and red indicates unstable solution.

x and the p_k 's are complex numbers. On the other hand, d_k^* is not needed, because d_k is a real number by its definition.

First we define a function $\mathcal{F} : \mathbb{C}^{3L+2} \rightarrow \mathbb{C}^{3L+2}$ such that

$$\begin{aligned}
\mathcal{F}_1(x, x^*, p_1, p_1^*, d_1, \dots, p_L, p_L^*, d_L) &= \dot{x} \\
\mathcal{F}_2(x, x^*, p_1, p_1^*, d_1, \dots, p_L, p_L^*, d_L) &= \dot{x}^* \\
\mathcal{F}_3(x, x^*, p_1, p_1^*, d_1, \dots, p_L, p_L^*, d_L) &= \dot{p}_1 \\
\mathcal{F}_4(x, x^*, p_1, p_1^*, d_1, \dots, p_L, p_L^*, d_L) &= \dot{p}_1^* \\
\mathcal{F}_5(x, x^*, p_1, p_1^*, d_1, \dots, p_L, p_L^*, d_L) &= \dot{d}_1 \\
&\vdots \\
\mathcal{F}_{3L}(x, x^*, p_1, p_1^*, d_1, \dots, p_L, p_L^*, d_L) &= \dot{p}_L \\
\mathcal{F}_{3L+1}(x, x^*, p_1, p_1^*, d_1, \dots, p_L, p_L^*, d_L) &= \dot{p}_L^* \\
\mathcal{F}_{3L+2}(x, x^*, p_1, p_1^*, d_1, \dots, p_L, p_L^*, d_L) &= \dot{d}_L
\end{aligned}$$

where $\{\dot{x}, \dot{p}_k, \dot{d}_k\}$ are defined in Eqs. 4.17. Then we derive the Jacobian matrix \mathcal{J} of \mathcal{F} , such that $\mathcal{J}_{ij} \equiv \partial \mathcal{F}_i / \partial v_j$.

The eigenvalues of the Jacobian matrix \mathcal{J} determines the stability of the steady state point. If signs of all the eigenvalues are negative, then the point is stable. If any of the eigenvalues is a positive number then we know that the point is an unstable steady state solution. Given a finite size of L , such calculations can be done very quickly on a laptop computer. The color code for the stability used in Figure 4.4 has been calculated this way, for example.

4.3.3 Effective Atom Numbers

When we drive the system with a very weak beam power, x (\propto a square-root of the intracavity photon number) will be small, and the steady state solution 4.22 can be

approximated by

$$y = x \left\{ (1 + i\Theta) + \frac{2C_0}{1 + i\Delta} \sum_{k=1}^L n_k s_k^2 \right\} \quad (4.25)$$

or

$$y = x \left\{ (1 + i\Theta) + \frac{2C_0 N_{\text{eff}}}{1 + i\Delta} \right\}, \quad (4.26)$$

where we defined the effective atom number N_{eff} as

$$N_{\text{eff}} \equiv \sum_{k=1}^L n_k s_k^2 = \sum_{j=1}^N c_j^2. \quad (4.27)$$

N_{eff} is a conceptually more useful quantity than N , the total number of atoms, for characterizing the system. Regardless of how the atoms are distributed inside the cavity and how the coupling strengths vary among those participating atoms, Equation 4.26 tells us that every system with the same N_{eff} would show the same input-output relationships under the weak drive. Even at higher input powers, the system with a wide spectrum of atom coupling strengths would show a behavior that closely follow that of an ideal system where N_{eff} atoms are all located identically at the center of the cavity (as studied in [19]).

The parameter N_{eff} also plays an important role in the actual experiment. After the atoms are loaded into the cavity mode by the use of the atom chip, right before and/or after some type of main measurement we usually make a secondary measurement solely to find the value of N_{eff} . We can add this extra measurement step without perturbing the atoms used in the main measurement process, because the N_{eff} measurement requires only 1 nW or less of the input power. The frequency of this input beam is constantly swept during the time, allowing us to record the transmission curve of the atom cloud-cavity system. The curve has two well-distinguished peaks,

which is well predicted by the values of x^2/y^2 (\propto output/input) from the Eq. 4.26 as a function of ω_{laser} . These two peaks can be understood as the normal mode splitting between the atom resonance frequency and the cavity resonance frequency. By fitting the measured transmission curve to the theory, we find N_{eff} .

The value of N_{eff} measured by this secondary measurement process is more reliable than what can be inferred by the main measurement result. One of the largest uncertainties in the experiment comes from the scaling factor that we use to convert the homodyne/heterodyne signals from the photodetectors to the real units of the optical power. On the other hand, the transmission curve x^2/y^2 from the Eq. 4.26 has no x or y terms in it, indicating that the locations of the normal mode peaks are independent of the input power. Therefore, we do not need to know the exact power of the input beam when we determine N_{eff} from the transmission curve (as long as it is small enough so that the approximation to Eq. 4.25 holds), making the measured value of N_{eff} immune to the possible errors associated with the signal-to-optical power scaling factor.

4.4 Modifications for nanophotonics applications

The general paradigm of two-level emitters coupled to an isolated optical mode can be realized in a wide range of gas, liquid, and solid-state systems with various types of optical cavities. For the case of our experiment, it was a Fabry-Perot resonator with an ensemble of laser-cooled ^{133}Cs atoms. The same principle can be implemented using ensembles of quantum dots [63] or vacancy centers [64] coupled to a nanophotonic cavity such as a microdisk [65, 66] or a photonic bandgap cavity [48, 67]. Such solid-state devices have much greater scalability and robustness than what can possibly

be achieved by the atom-cavity systems in ultra-high vacuum, making them good candidates for future photonic computational devices [22, 68]. Commercialization of these devices can also benefit from the infrastructures of the semiconductor industries as well. Moreover, unlike the MHz scale dynamics we observed in our experiment (Chapter 5), the intrinsic time scales in those nanophotonic devices will be much faster, in GHz scale [21], due to their small physical dimensions. The GHz scale dynamics will allow them to be more compatible with current computation devices.

Spectral distribution

In order to apply the Maxwell-Bloch equations developed so far to the possible solid-state implementations, some adjustments are needed to correctly model the peculiarities of the nanophotonic systems. First of all, the resonance frequencies of the virtual atoms in solid-state devices are expected to have a wide spectral distribution. These virtual atoms may still be approximated as two level systems but it will be a big challenge making them all share the same resonance frequency. This is in contrast to the case of the real atoms used in our experiment, where all the atoms are physically identical, allowing them all to have the same resonance frequency as long as the background magnetic field is uniform. To include the spectral distribution to the Maxwell Bloch equations, Eq. 4.12b must be changed to

$$\langle \dot{\sigma}_{-}^j \rangle = -(\gamma_{\perp} + i\Delta_{A,j}) \langle \sigma_{-}^j \rangle + g_j \langle \mathbf{a} \rangle \langle \sigma_z^j \rangle \quad (4.28)$$

where $\Delta_{A,j}$ is the detuning for the j -th atom.

Dephasing

We considered only the radiative damping (with the rate γ_{\perp}) of the atoms in our model. Considering the extremely long coherence time (> 1 second) of the ground state atoms in the vacuum, it was not necessary to include the non-radiative damping in the master equation 4.2. However, in the solid state implementation the virtual atoms interact much strongly with their environments and they can transfer their excitation energy to their neighboring atoms or through phonons and lose their energy incoherently. This leads to an extra dephasing process. We denote such non-radiative dephasing rate by γ_{nr} and we modify the the unconditional master equation to

$$\begin{aligned} \frac{d}{dt}\rho = & -i[\mathbf{H}, \rho] + \frac{\gamma_{\parallel}}{2} (2\sigma\rho\sigma^{\dagger} - \sigma^{\dagger}\sigma\rho - \rho\sigma^{\dagger}\sigma) + \kappa (2\mathbf{a}\rho\mathbf{a}^{\dagger} - \mathbf{a}^{\dagger}\mathbf{a}\rho - \rho\mathbf{a}^{\dagger}\mathbf{a}) \\ & + \frac{\gamma_{nr}}{2} (\sigma_z\rho\sigma_z - \rho), \end{aligned} \quad (4.29a)$$

where we now use γ_{\parallel} to denote the spontaneous decay rate for the atoms, instead of $2\gamma_{\perp}$ ⁴. Refer to [19] to see how this changes the Maxwell-Bloch equations.

⁴In the presence of non-radiative decay, γ_{\perp} is now defined as $\gamma_{\perp} \equiv \gamma_{\parallel}/2 + \gamma_{nr}$.

Chapter 5

Optical Bistability

5.1 Experiment Overview

In our experiment, the atoms are first collected and cooled by a magneto-optical trap (MOT) formed on a gold-coated mirror surface [45]. This surface MOT is located outside the cavity, 2 cm away from the cavity axis, where the cooling beams are not obstructed by the cavity mirrors. The atoms are then further cooled to sub-Doppler temperature ($< 10 \mu\text{K}$) by polarization-gradient cooling, optically pumped to the $m_F=4$ state, and then loaded in the magnetic trap. The magnetic trap, formed by currents through copper wires (100 – 200 μm diameter) buried underneath the mirror surface, can be moved by shifting the currents [57], and is used as a conveyor to transport atoms (up to $\sim 10^6$) into the cavity mode. The number of atoms to be transported are adjusted by changing the initial MOT size.

At the end of the transport process (200 – 300 ms), the atoms are released from the magnetic trap and a large uniform magnetic field is applied in the direction of the cavity axis, which ensures that the quantization axes for the atoms and the field are

aligned in the presence of possible stray fields near the surface. The strong bias field also lifts the degeneracy among the Zeeman states (Zeeman shift: $2\pi \cdot 17$ MHz). Once the atoms are in position, the input laser is injected from one side of the Fabry-Perot cavity and drives the TEM₀₀ mode. The beam that leaks out from the other side of the cavity serves as the output of the system. Due to the loss at the cavity mirrors and the existence of two possible directions for the leakage, we estimate η , the probability of an intracavity photon making it to the output channel, to be 0.41. The output beam goes through an optical isolator and is then measured by a homodyne/heterodyne detection scheme [10].

The maximum atom-cavity coupling constant, the atomic polarization decay rate, and the measured cavity field decay rate are $(g_0, \gamma_\perp, \kappa)/2\pi = (0.72, 2.6, 2.3)$ MHz, respectively.

5.2 Input-Output Characteristics

In Figures. 5.1(a) and 5.1(b), we plot the measured input-output response of our atom-cavity system for two different parameter sets $\{N_{\text{eff}}, \Delta_C, \Delta_A\}$. For these measurements the input power was swept slowly (20 kHz) compared to the dynamical rates of the system $(g_j, \gamma_\perp, \kappa)$ to probe the adiabatic response. The data traces were recorded continuously over $\sim 50 \mu\text{s}$ as the input power was swept up and down.

For the parameters of Fig. 5.1(a) ($N_{\text{eff}} = 130, \Delta_C = \Delta_A = 0$) a soft thresholding behavior is observed with a corner near 6 nW input power. There is no significant difference between the up-sweep measurement result and the down-sweep one, confirming that the system we probed stays very close to the steady state and the data we measured contain little transient behavior.

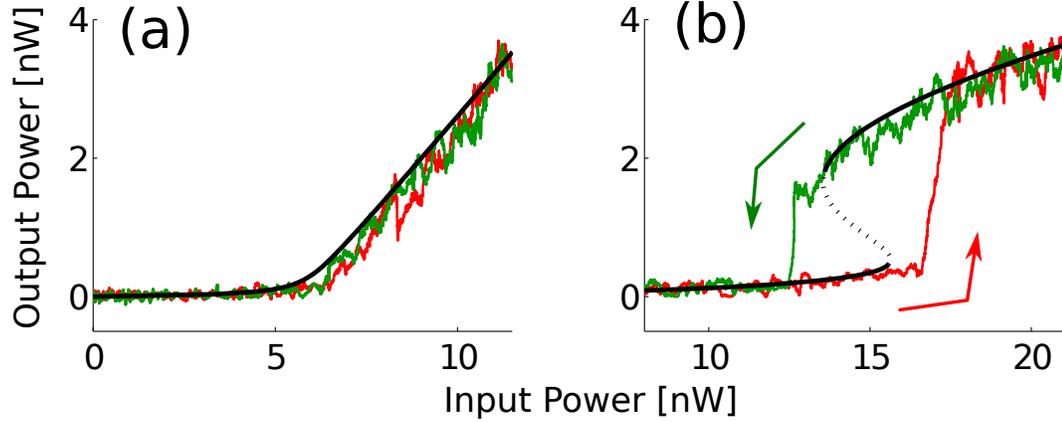


Figure 5.1: (a) Threshold behavior of the system with $N_{\text{eff}} = 130$, $\Delta_C = \Delta_A = 0$. The black line is the stable steady state solution derived from the model, and red (green) line is a typical data trace from the measurement when the input power is gradually increased (decreased). (b) Hysteretic behavior with $N_{\text{eff}} = 470$, $\Delta_C = \Delta_A = 2\pi \cdot 8$ MHz. Color scheme is the same as (a) except for the new black dotted line that indicates the unstable steady state solution.

On top of the measurement data drawn in red and green, we also plot in black corresponding steady state solution Eq. (4.23) of the Maxwell-Bloch Equation. The match between the theory and the measurement is pretty good. The theory curve was actually manually fitted by adjusting the value of N_{eff} so that the agreement with the measured data is somewhat optimized, but we also make sure that this fitted N_{eff} is not too different from the independently measured N_{eff} from the experiment.

To be more specific, N_{eff} was recorded (for each experimental run) by measuring the spectral response of the atom-cavity system using a very low intensity probe [29], as explain in section 4.3.3, right before and after the main measurements were made. We have found that the measured values of N_{eff} from this procedure differ from the fitted values by no more than 10% in all of our measurements. This holds not just for the thresholding behavior measurements we show in Fig. 5.1(a), but also for

all the other measurements I present in this thesis. The comparison of these two spectral curves, before and after the measurement, also ensures that N_{eff} does not vary significantly during the measurement period.

With more atoms coupled to the system, the nonlinear behavior of the system becomes more prominent than mere thresholding, and we can observe classic bistability and hysteresis. For Figure 5.1(b), we now used $N_{\text{eff}} = 470$, $\Delta_C = \Delta_A = 2\pi \cdot 8$ MHz, while we swept the input power to the system in the same manner as in Fig.5.1(a). Unlike the thresholding case, however, we find that the steady state output we get from up-sweeping and down-sweeping produce different input-output plots.

We measure that the coexisting stable equilibria exist from roughly 14 to 15.5 nW input power. When we slowly increase input power from below the bistable region, the system will stay in the lower branch until the power reaches 15.5nW, at which point the output suddenly jumps up. Similarly, when we bring down the input power from above the bistable region, the system stays in the upper branch, until it makes a down jump at 14nW. When compared with the theory plot, we clearly observe that this hysteresis behavior is exactly what we should expect from the ‘S’-shaped steady state solutions. Jumps happen when the steady states disappear at bifurcation points where stable and unstable steady state solutions meet.

5.3 All-Optical SR Latch

After confirming the parameters for simple bistable input-output behavior in our atom-cavity system, we performed a proof-of-principle experiment to demonstrate an all-optical set-reset (SR) latch, which could for example be used as an optical memory bit. Typical data is shown in Fig. 5.2. Here we envision a device configuration in

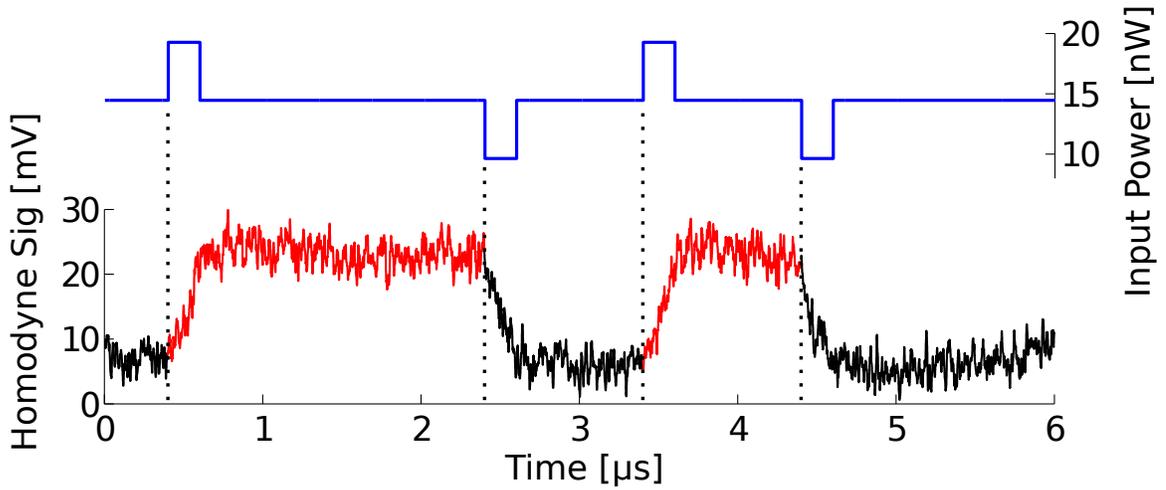


Figure 5.2: Experimental demonstration of latching using bistable system parameters of Fig. 5.1(b). With the input power constant at 14.5nW the system maintains its current state. When a pulse is applied the system state is set high or low according to the pulse type (positive/negative).

which the set/reset control beam shares an optical input channel with the bias power beam (for example they could be combined by a beam splitter) so that they interfere constructively or destructively depending on the phase of the control beam. With this picture in mind, we demonstrated the set/reset control by superposing positive and negative pulses onto the laser beam that drives the atom-cavity system (upper blue trace of Fig. 5.2). A positive pulse switches the response to the upper branch of the hysteresis curve, where it latches until a negative pulse resets it back to the lower branch. Our system required 15 nW optical bias power and the energy used for each control pulse was roughly $5 \text{ nW} \times 0.2 \mu\text{s} = 1 \text{ fJ}$.

We note that the width of the control pulses used in our experiment corresponds to the inverse cavity decay rate κ^{-1} . Our switching time is thus longer than in nanoresonator-based devices, which generally have much larger κ due to their small size. However, in terms of the switching energy, ours appears to be the minimum

energy demonstrated in an all-optical bistable device [14]. The switching energy potentially could be lowered even more as the energy scale in this experiment is still much larger than the optical shot noise (quantum fluctuation) limit [10].

Chapter 6

Limit Cycle

6.1 Instability Leading to a Limit Cycle

With an appropriate choice of external parameters N_{eff} , Δ_C and Δ_A guided by theory, we also probed the atom-cavity dynamics in an unstable regime and observed “self-pulsing” behavior [2, 38] in which the global attractor is a limit cycle.

Finding a parameter set that leads to an instability was done somewhat by trial and error. The steady state solutions over a wide range of input power can be quickly calculated, along with their stabilities. One example of such calculation is shown in Figure 6.1. When we get a input region in which there is only unstable steady state solution, as shown as the red segment in the figure, we can assume that the system will be repelled from the steady state solution at these input powers and will end up in a stable orbit in the phase space called a limit cycle ¹. Numerical simulation of this behavior can be easily done. The right plot of Fig. 6.1 shows the simulated

¹Other possibility is a chaotic trajectory in the phase space [34, 35, 69, 70], which I have not encountered in my system.

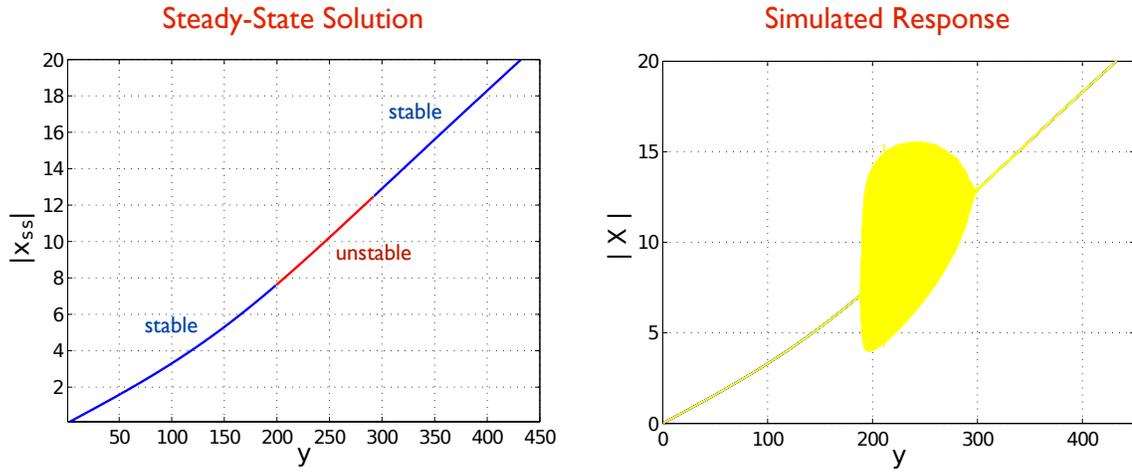


Figure 6.1: [Left] Steady state solution for Maxwell-Bloch equations 4.8 with $(C, \Theta, \Delta) = (47, -20, 7.7)$. The stability for each point on the plot is also calculated and the plot is painted with blue (stable) or red (unstable) accordingly. [RIGHT] With the same parameters, Equations 4.8 is integrated over time to simulate the evolution of the states. During the simulation, the input power (y) was zero at $t = 0$ and was slowly swept up.

cavity field amplitude ($|x|$) as the input drive was sweeping. The blob in the center is an oscillating trace with a well-defined frequency.

In Fig. 6.2 we show the experimental measurement of the limit cycle behavior. We can see the transition of the system between a stable equilibrium point and a limit cycle as we sweep the driving input power. When the input power is large, the power of the output beam oscillates significantly at a frequency in the range of $\sim 21 - 24$ MHz. Note that more bandwidth is retained in the homodyne signal displayed here than in Fig. 5.1 in order to preserve the limit cycle oscillations.

We also tested the system's response to a sudden change in the input power. In Fig. 6.3, we switch the input power between 100%, 50% and 0%, where the limit cycle is possible only when the power level is 100%. The system jumps to a limit cycle or

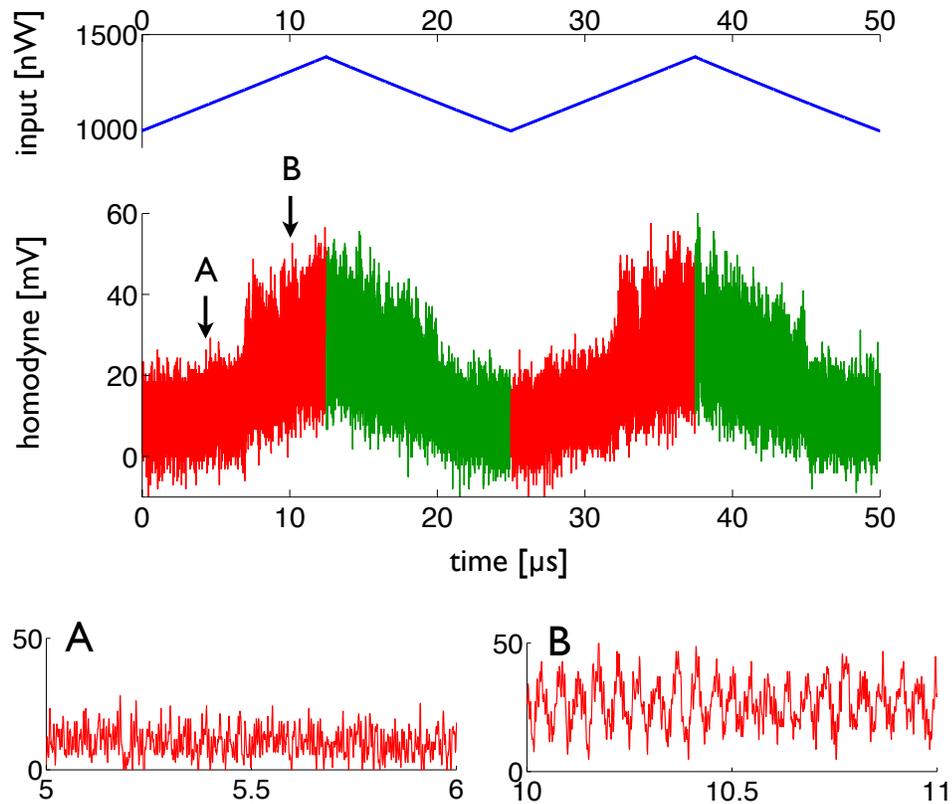


Figure 6.2: Experimental parameters $N_{\text{eff}} = 2100$, $\Delta_C/2\pi = -20$ MHz, $\Delta_A/2\pi = 5$ MHz are used. Homodyne signal generated by the output beam as the input power is slowly swept up and down over a Hopf bifurcation point. Zoomed-in views for before (A) and after (B) the bifurcation point are displayed on the bottom.

jumps back to a stable equilibrium point very fast in response to the change in the input power.

6.2 Sub-Critical Hopf Bifurcation

In Figure 6.4, we show both the calculated steady state solutions and the simulated max/min output from the limit cycle for the system with the same parameters used in Fig. 6.2. An interesting behavior of the system displayed in this figure is the hysteresis

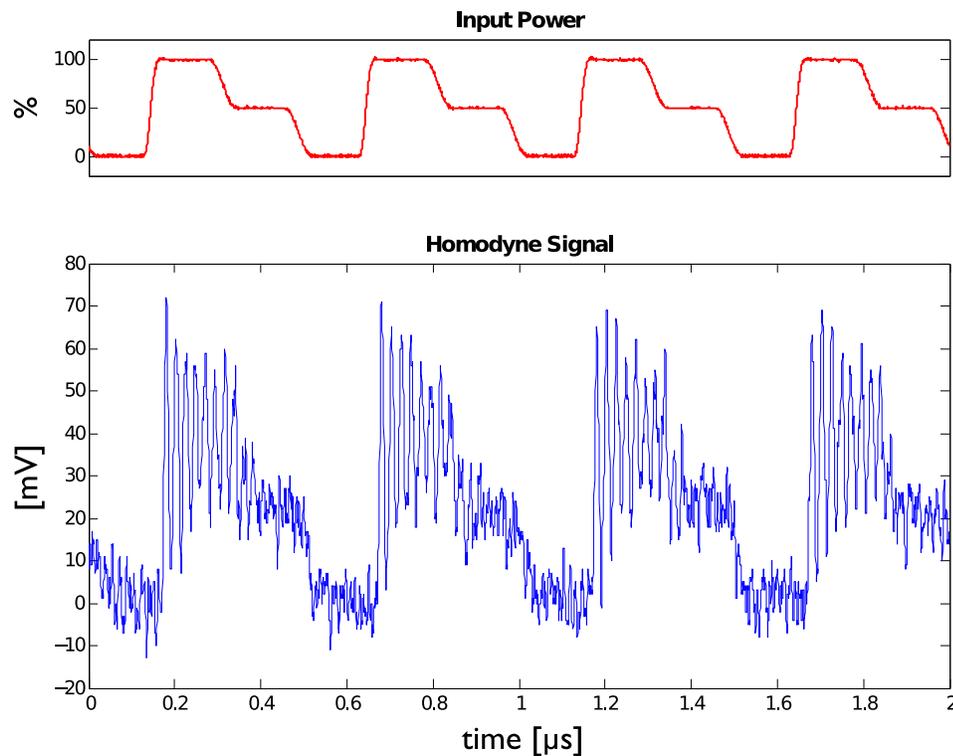


Figure 6.3: Response of the system to sudden jumps in the input power. Only when the full power (100%) is used, the system makes the transition to a limit cycle.

around the sub-critical Hopf bifurcation point². Our model predicts a range of input powers ($\sim 1150 - 1200$ nW) for which a stable equilibrium point and a stable limit cycle coexist. Starting from low input, as the system is driven past the sub-critical Hopf bifurcation point (~ 1200 nW), we expect that the equilibrium point becomes unstable and the system jumps to the the limit cycle. On the other hand, with the system oscillating, the input power can be lowered back down and the limit cycle remains until another bifurcation occurs at a lower input power (~ 1150 nW) than the sub-critical Hopf bifurcation point.

²I will not explain what Hopf bifurcations are in this thesis. Readers can look up any nonlinear dynamics textbooks, such as [71].

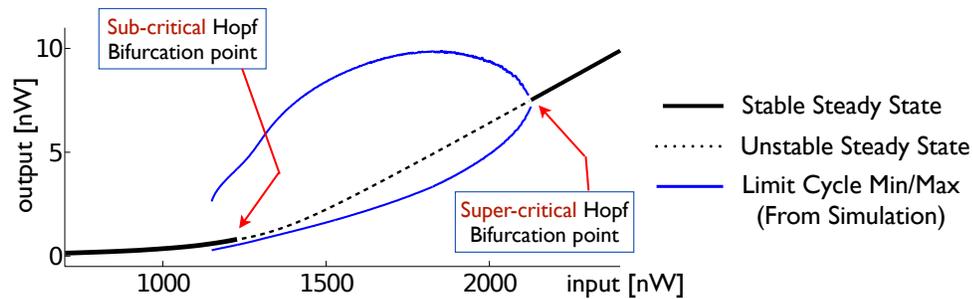


Figure 6.4: Theoretical prediction of the input-output characteristics of our system that shows instability. The solid and the dotted black lines represent the stable and the unstable steady state solutions, as in Fig. 5.1B. We also simulated the system with a slowly varying input, as in Fig. 6.1, and searched for input ranges where the limit cycle happens. The maximum (minimum) power of the oscillating output for each input power is recorded, and is plotted here as the upper (lower) blue trace.

This behavior is qualitatively different from what we expect to see near the super-critical Hopf bifurcation point. Simulation predicts that the amplitude of the limit cycle gradually converge to zero as the system approaches the super-critical Hopf bifurcation point from below, at which point a stable equilibrium point defines the steady-state. When the input is swept in the opposite direction, no hysteresis is found. The steady state system starts to develop oscillation with near zero amplitude, without any sudden jumps.

The comparison between the theory and the measurement near the subcritical point is shown in Fig. 6.5. On the top, we plot the frequencies of the both simulated and measured limit cycles. The simulation result is plotted in brown, while the measurements are denoted by red (up-sweep) and green (down-sweep) square marks. The match between the theory and the experiment seem reasonable, and the measurement correctly shows the increase of the limit cycle frequency as the input power is lowered and the amplitude of the limit cycle is reduced.

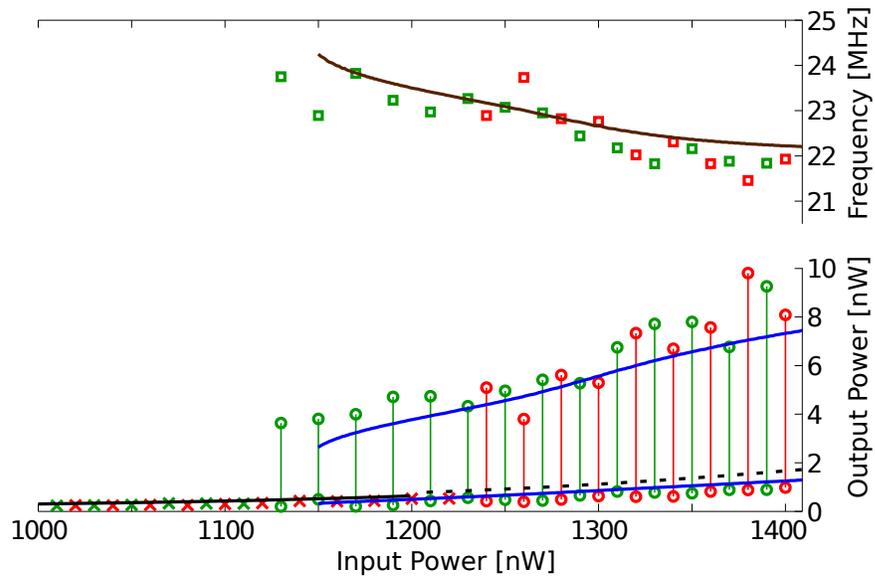


Figure 6.5: Measured and predicted output data for experimental parameters $N_{\text{eff}} = 2100$, $\Delta_C/2\pi = -20$ MHz, $\Delta_A/2\pi = 5$ MHz. Top plot is the limit cycle frequency, and the bottom plot is the output power from the system. Refer to the main text for more information.

On the bottom of the Figure 6.5, the theoretical prediction of the input-output characteristics are plotted with black and blue lines, same as in Fig. 6.4. The measured output is plotted in red (up-sweep) and green (down-sweep), using the following rules: When no oscillation is detected, the mean output power is recorded by a cross-mark. When the oscillation is detected, a sinusoid is fitted to the measured signal and its max/min optical power is denoted by a pair of connected circles. The hysteric behavior is clearly observed, although there is some discrepancy with the theory on exact location of where the jumps happen.

6.3 All-Optical DC-RF Latch

Demonstration of the DC-RF Latch

The hysteresis around the sub-critical Hopf bifurcation can be used to make a latch of rather unique characteristics, which switches between dc and modulated signal formats. Optical dc-rf latching devices may find use in optical communication systems where signal up-conversion is a common practice to avoid the high level of the background noise at low frequencies. Even though the response time of our cavity QED system is rather slow and thus may look not so attractive in the practical sense, we note that solid-state implementations of this type of cavity nonlinear optical device, for example using ensembles of quantum dots [63] or vacancy centers [72], could achieve much higher values of $(g_0, \gamma_{\perp}, \kappa)$ and therefore much faster response, and much higher limit cycle frequencies, potentially in the GHz range.

Using the same system shown in Fig. 6.2, but this time by applying high and low pulses in an otherwise constant input power, we demonstrate such latching operation between a stable equilibrium point and a limit cycle. The result is shown in Figure 6.6. When a positive pulse is applied, the system switches its state to RF and remains in the limit cycle. When a negative pulse is applied, it goes back to a low-power steady state. In the bottom figure, we show a zoomed-in view of the homodyne signal from $3 \mu\text{s}$ to $4 \mu\text{s}$.

Figure 6.7 compares the measured output field quadratures for both rf and dc state of the latch with the theoretical prediction on the optical phase plane. Here, it is a little difficult to see the good match because of the noise in the quadrature measurements, but the distinction between the dc and rf output is clear. The field quadrature data measured for rf state even shows the ellipticity of the limit cycle on

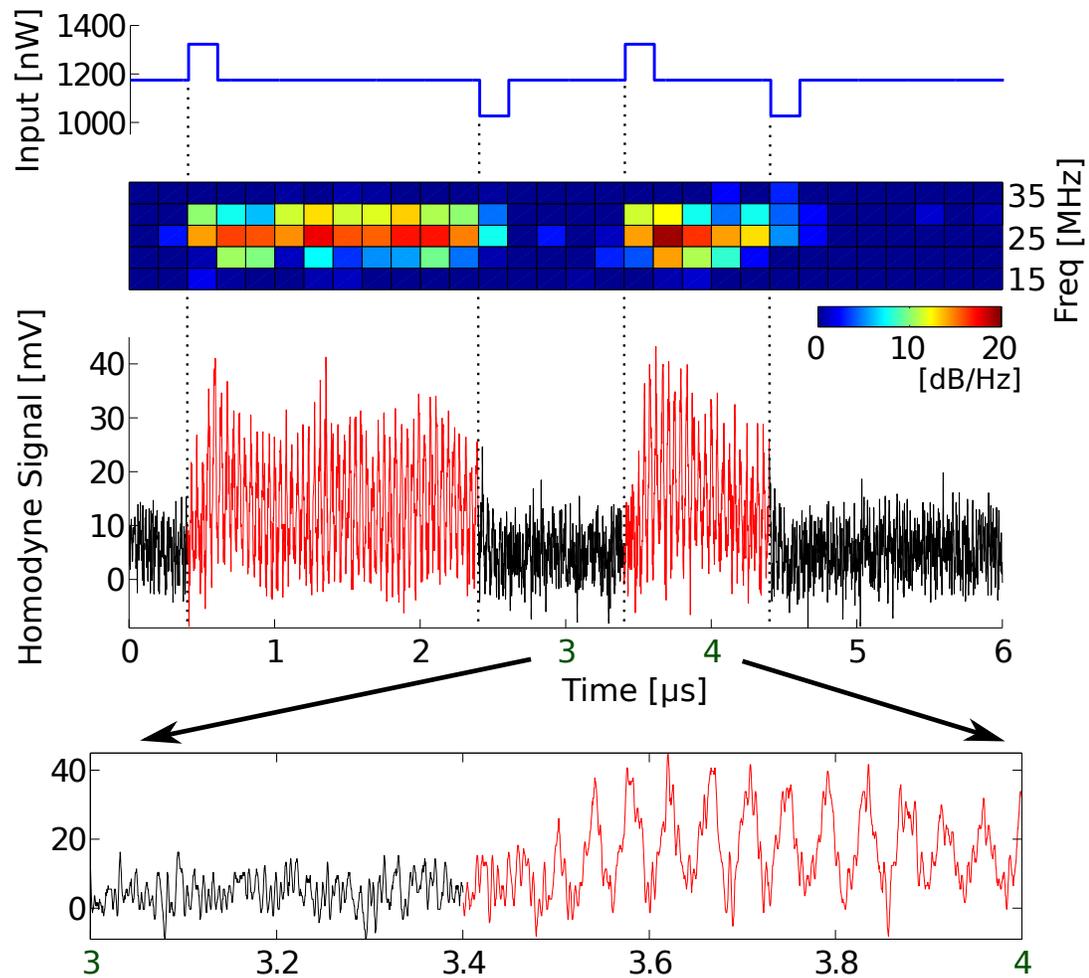


Figure 6.6: A demonstration of all-optical DC-RF latch. The spectrogram shown in the middle of the figure is generated using $0.2 \mu\text{s}$ segments of the homodyne data, which contain roughly 5 oscillation cycles. 0 dB power level was arbitrary chosen.

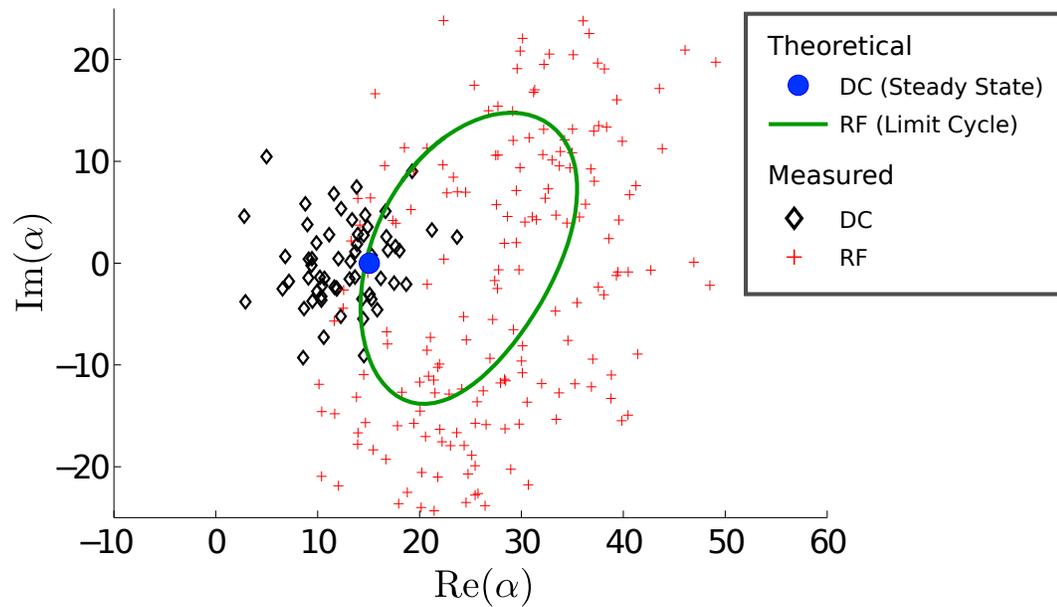


Figure 6.7: The measured dc and rf states of the latch are compared with theoretical predictions. Here, α is the complex amplitude of the coherent cavity field state. The measured data for the rf state plotted in the figure represent 25 cycles of oscillation, roughly 7 data points for each cycle.

the phase plane correctly. By the way, the noise we see in this figure is a pure quantum shot noise. The width of the data spread matches the expected level of quantum noise on the heterodyne signal output, accounting for the quantum efficiency of the detector ($\sim 20\%$) and the optical losses.

How the Phase-Space Plot (Figure 6.7) Was Made

In creating Fig. 6.7, because we had to map out the oscillation of the rf state output on the phase diagram, not a steady state output, we faced two difficulties. First, we could not improve the signal-to-noise ratio as much as we wanted, because we could not integrate the quadrature measurements for longer than a fraction of a limit cycle period. Otherwise, the movement of the system around the limit cycle cannot be

resolved. This is why the data we show had to contain such large spread.

The second problem was the difficulty in measuring the complex quadrature values itself. Normally, in order to measure the complex amplitude of an optical field, the heterodyne detection is used, where this field is interfered with a stable, frequency-detuned local oscillator. The information carried by the measured signal oscillates between that of the amplitude quadrature ($\text{Re}\{\alpha\}$) and that of the phase quadrature ($\text{Im}\{\alpha\}$) at the (fast) detuning frequency. Each value of the field quadrature is then extracted through rf demodulation of the heterodyne signal. In our case, however, such a scheme was not applicable because the intracavity field state we wanted to measure oscillates at a fast rate ($\sim 23\text{MHz}$), comparable to the bandwidth of the detector (80MHz).

For this reason, we used a slow heterodyne frequency (2MHz) instead, so that we effectively measured only one quadrature value while the intracavity state underwent several limit cycle oscillation. Assuming that the known limit cycle frequency remains constant, we define a set of points in time, spaced by the measured period of the limit cycle, such that the state of the system at each time is supposed to be on the same spot of the limit cycle. From this set of points, we select two points whose separation is close to the quarter period of the heterodyne measurement (125 ns). Measurements made at these two points in time give us a pair of complementary quadrature values, which we then combine to create one complex amplitude data point. We repeat this process over ~ 25 limit cycle oscillations to create the complex phasor portrait displayed above.

Switching Energy for the DC-RF Latch

We find that the energy stored in the optical dc-rf latch in Fig. 6.6 is comparable to that of the elementary latch in Fig. 5.2, as in both cases the cavity contains about 500-1000 photons (100-200 attojoule) on average in the high energy state and much fewer when in the low energy state. However, the energy required for switching between the two states was more than an order of magnitude higher for the dc-rf latch (~ 30 fJ). This is mostly due to the mismatch between the input laser frequency and the resonances of the system.

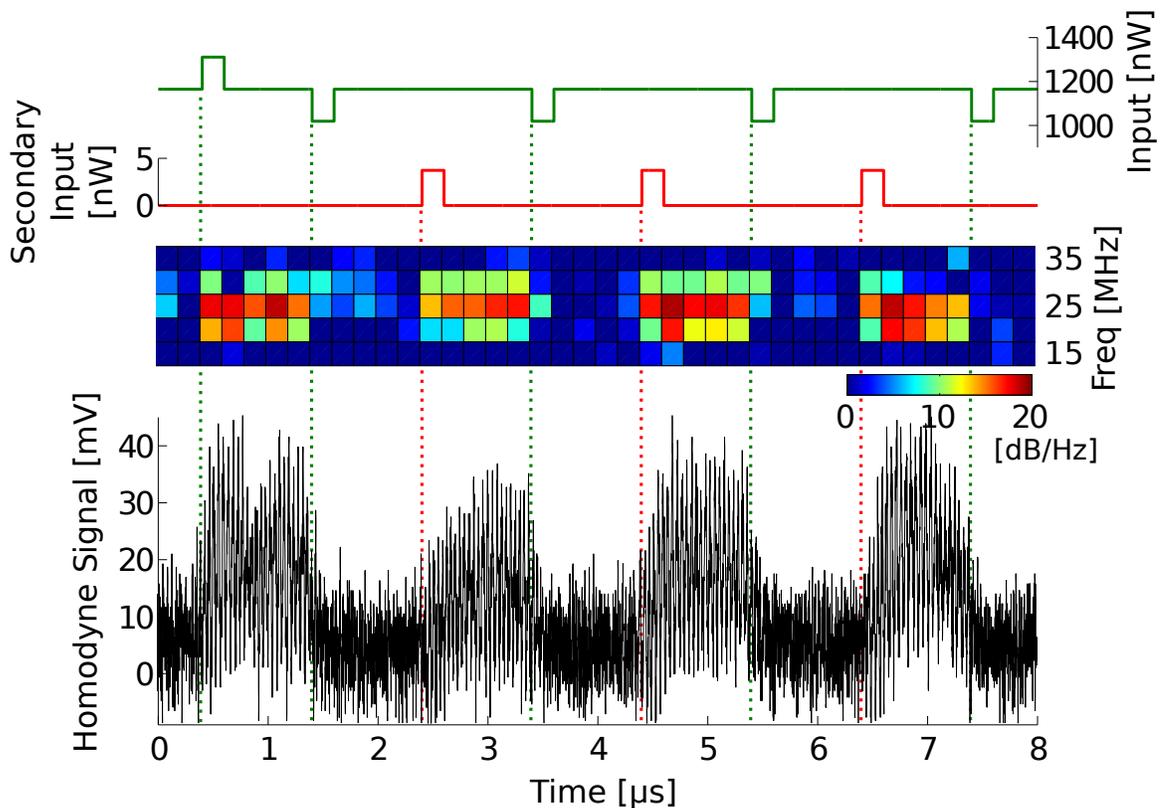


Figure 6.8: Demonstration of the dc-rf latch switch-on operation using a sub-femtojoule secondary input pulse.

We can alleviate the frequency mismatch problem by using a secondary near-resonant beam. In Figure 6.8, we added a second input beam which has frequency ω_{las2} where $\omega_{\text{las2}} - \omega_{\text{las}} = -2\pi \cdot 23\text{MHz}$ and $\omega_{\text{las2}} - \omega_{\text{cav}} = -2\pi \cdot 3\text{MHz}$. The beating created by the two input beams matches the expected frequency of the limit cycle.

During the 0 - 2 μsec , the dc-rf latch was switched on and off by adding positive and negative pulses to the primary beam (shown in green), similar to what was demonstrated in Fig. 6.6. In contrast, from 2 to 8 μsec rf latching was initiated by applying pulses of the secondary beam (red) with energy $4\text{nW} \times 0.2\mu\text{s} = 0.8\text{fJ}$ for each pulse. We have observed that at this energy level the pulses are able to switch the latch from the dc- to the rf-state reliably. Negative pulses of the primary beam were still used to reset the latch back to the dc-state and have it ready for subsequent switch-on pulses.

We also note that the use of pulse shaping on the control pulse [73] could potentially reduce the switching energy of both the elementary and the dc-rf latch further.

Chapter 7

Optical Nonlinear Amplification

In this chapter, I present the results of my experiment on the phase-insensitive all-optical nonlinear amplification. To our knowledge, these results represent the first demonstration of the optical amplification using the small-signal amplification scheme described in 7.1, or any all-optical nonlinear amplification at that. Dr. Soh from Sandia National Laboratories is currently continuing the experiment and probing many aspects of this phenomenon that are not yet included in this thesis. Nikolas Tezak and Ryan Hamerly in our group are working on the theoretical side of this project as well. Hopefully all of our efforts can be combined in the near future so that we can produce a coherent, more complete analysis of this new phenomenon we observed. In the mean time, I will focus in this thesis on presenting the first few series of data in which we tested the two most basic characteristics of the amplification, namely, the amplification gain and the noise.

7.1 Small-Signal Amplification Using Super-Critical Hopf Bifurcation

As was pointed out in [74], when a system can exhibit instability that leads to a limit cycle, such a system near its bifurcation point can amplify a small periodic signal that perturbs the system at the expected limit cycle frequency. We have already shown in the previous chapter that our atom cloud-cavity system, under a certain set of parameters, shows the limit cycle behavior. A direct application of the above amplification scheme would mean that we bring the system close to the bifurcation point and modulate the input beam at the limit cycle frequency. Then we are expected to see a larger modulation on the output beam. While such a direct approach is slightly different from what we actually implement in our experiment, it conveys the general idea how our amplification scheme works.

In order to explain the subtle difference between our experiment and the straightforward amplification scheme described above, we first point out the asymmetry found in the power spectrum of the output beam when our system undergoes the limit cycle behavior. Figure 7.1 shows three different output power spectra resulting from three different input powers selected near the super-critical Hopf bifurcation point. The case A and B represent the instability cases with limit cycles, which create the sidebands next to the carrier waves in the power spectra. The carrier is located at the input beam frequency, and the distance between the carrier and the sideband is the limit cycle frequency, which is usually in tens of MHz. If the modulation on the output created by the limit cycle had been pure amplitude modulation (AM) or pure frequency modulation (FM), the two sidebands would have the equal power. But as

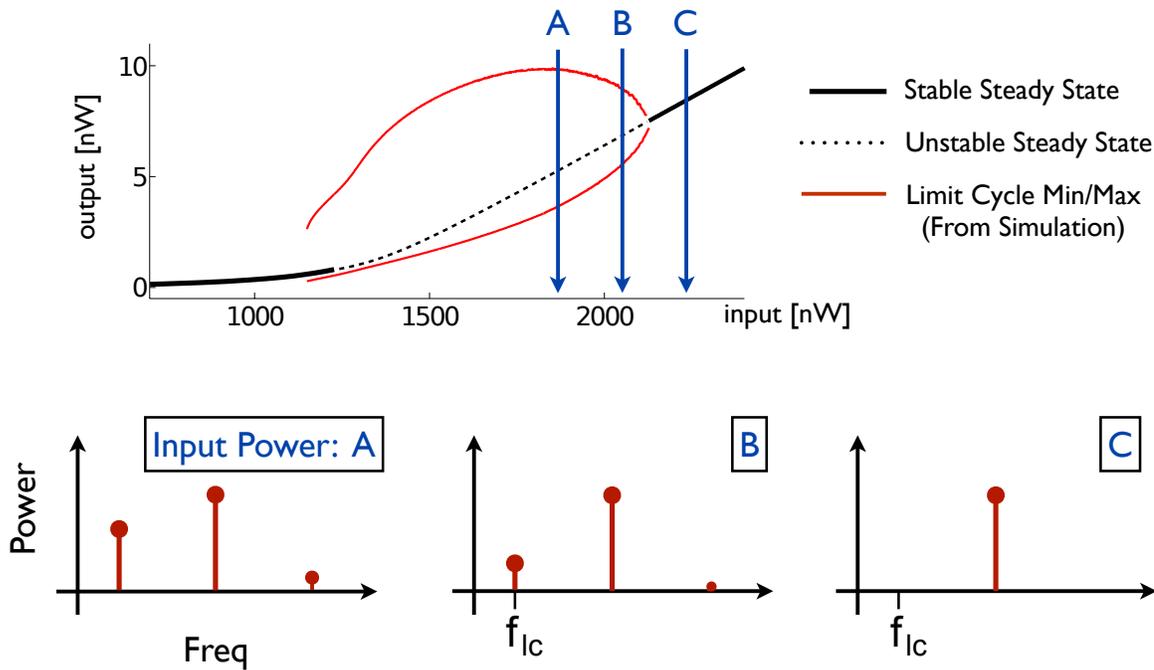


Figure 7.1: On the top shows the same input-output characteristics curve of Fig. 6.4. We indicated three different power levels of the input beam and marked them by A, B, and C. On the bottom, we show how the power spectra of the output beams look like at each of those three input power levels.

was already pointed out in Fig. 6.7 the limit cycle modulation is roughly circular in the phase diagram, which leads to one sideband being much stronger than the other. This explains the asymmetry in the power spectra.

Ignoring the smaller sideband of the two (and other possible higher order sidebands that are negligibly small), all the limit cycle does is adding a second peak in the power spectrum. In other words, we can think of it as simply producing a second beam with an optical frequency f_{lc} and adding it on top of the transmitted input beam, where f_{lc} is the frequency of the major sideband as marked in Fig. 7.1. The fluctuation in the output intensity, in this sense, is just a beating effect between the two beams.

Now, if the input power is increased past the super-critical Hopf bifurcation point,

such as the case C in Fig. 7.1, the system becomes stable and there is no more limit cycle behavior. The output beam is now back to being monochromatic and the power spectrum shows only a single peak. However, some eigenvalues of the Jacobian matrix (described in section 4.3.2) at that steady state would have negative real parts but still very close to zero. This means that the system is prone to become unstable given a small perturbation and may start lasing at the frequency f_{lc} again.

In our amplification experiment, we superpose a very weak beam of the frequency f_{lc} , which we call either a seeder or a signal beam, on top of the strong input beam, which we now call the pump. The pump power is chosen so that the system is right above the super-critical Hopf bifurcation point. By the nonlinear dynamics described above, the output contains not just the transmitted portion of the seeder beam but also an amplified power at the frequency of the seeder, as well as the well-defined phase relationship with it.

It is interesting to think about how this amplification actually works. The amplified output is not coming from the stimulated emission from the intracavity atoms as in the case of laser amplifiers, because the intracavity atoms do not provide a population-inverted gain medium. The optical power is converted from the pump to the seeding laser through the internal dynamics of the system, which is regulated by the externally injected seeder. The amplification is phase-insensitive, not because the system can amplify the seeding laser of an arbitrary phase, but because the system adapts its phase to the phase of the seeding laser. More specifically, the phase (and therefore also the frequency) of the limit cycle dynamics is synced to the injected beam, allowing the system to generate the beam with copied characteristics of the seeder but with a larger power. In this sense, our amplification system is more like

a phase-locking optical oscillator, except that the oscillator's output power is not constant but generally proportional to the power of the seeder.

7.2 Measured Amplification Gain

Before we carry out the amplification experiment, we first need to find the exact location of the super-critical Hopf bifurcation point. In the case of the sub-critical Hopf bifurcation point, finding where it occurs is very obvious because the system makes a noticeable jump at the bifurcation point during the slow sweep of the input power. There is no such clear-cut indication of the super-critical Hopf bifurcation point, on the other hand, since the change from the instability to the stability, or vice versa, is supposed to be continuous at the bifurcation point.

In Figure 7.2, we display the measured output power spectrum at different levels of input power (no seeder, pump only). At the pump power of 950 nW, we see a strong peak originating from the limit cycle. As the input power is increased (950 nW to 1100, 1300, and 1400 nW), the limit cycle amplitude becomes weaker, as is expected for a system that approaches the super-critical Hopf bifurcation point. We also expected that the peak would completely disappear as the input power is raised past the bifurcation point, but this was not what we observed. At the input power 1400 nW and above, a small peak persistently remains, making it impossible to pinpoint, or even properly define, the bifurcation point.¹ Nevertheless, seeing that the size of the limit cycle at 1400 nW input is already as small as it can get, we decided to use this power level as the pump power in the following amplification experiments.

¹A more advanced model that incorporates the noise of the system (such as the shot noise on the input beam) into the Maxwell-Bloch equations can correctly predict such non-vanishing small limit cycle behavior. Ask Ryan Hamerly for more details.

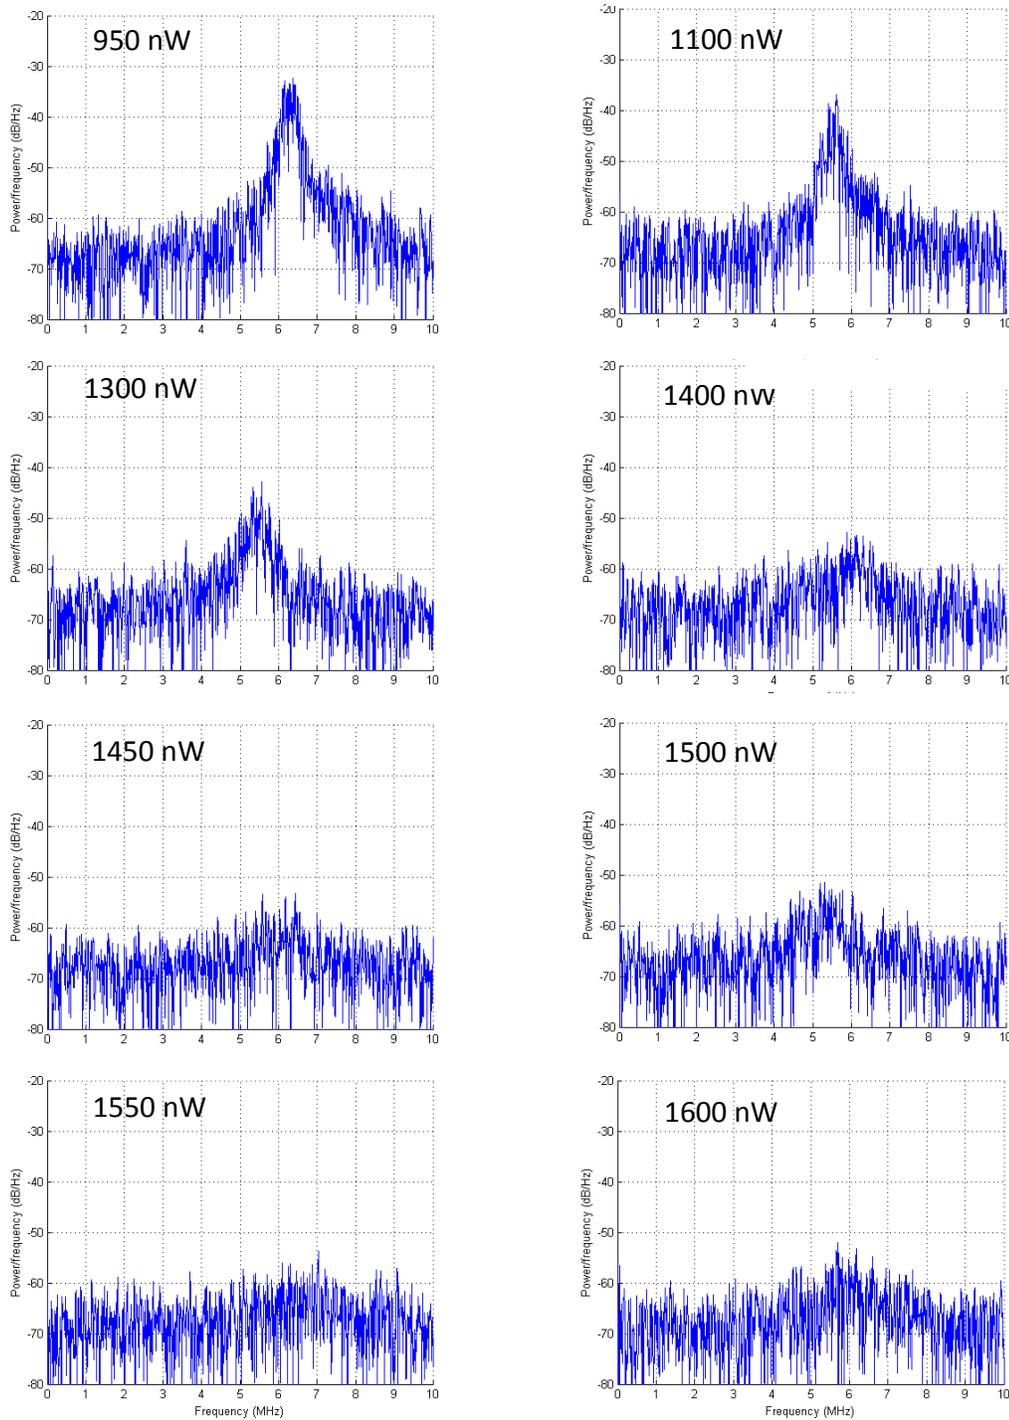


Figure 7.2: Power spectrum of the limit cycle at different input (pump) powers. Used parameters are $N_{\text{eff}} = 2000$, $\Delta_C = -20$ MHz, and $\Delta_A = 5$ MHz. The x-axes in these plots represent detunings from the frequency of the local oscillator beam used for the heterodyne detection, -22 MHz red-detuned from the atom resonance. These displayed spectra are simply FFT results of the heterodyne measurement signal.

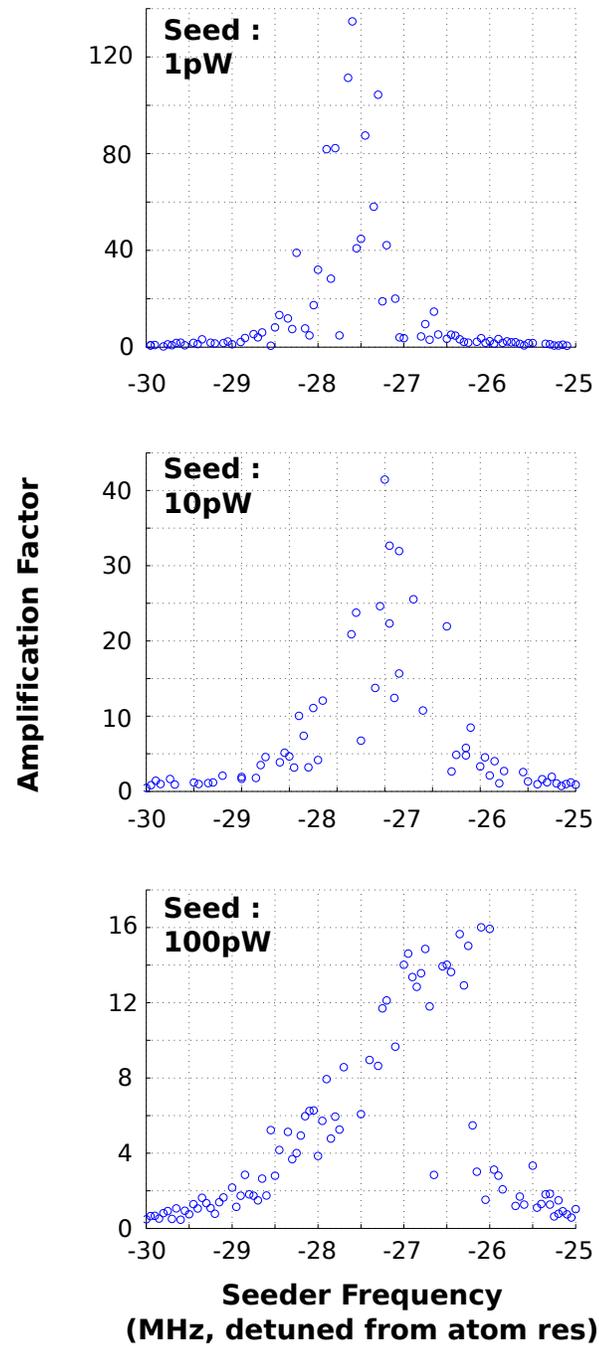


Figure 7.3: Gain curve of the demonstrated optical amplifier. Used parameters are the same as Fig. 7.2 with pump power set at 1400 nW.

Figure 7.3 shows the gain curves of the optical amplifier with pump power of 1400 nW. Each point in the plot comes from a roughly 100 μ s measurement result, each measured in a separate run with freshly collected, cooled and transported atoms. Before and after each measurement, we record the transmission curve of the system which is later used to determine the effective atom number for that particular run (See section 4.3.3). Only the data tagged with atom numbers between 1900 and 2100 is plotted in Figure 7.3 and the others are discarded.

The output field contains both the transmitted pump beam and the amplified seeder beam. The raw heterodyne measurement data recorded by an oscilloscope is first treated to filter out the pump beam signal which is much stronger than the amplified seeder signal.² After this, what is left gets mixed with the sine and cosine function of the seeder signal frequency, giving us the two complementary quadrature values based on an arbitrary phase reference. Combining the two quadrature values, we finally obtain the ‘real’ amplitude quadrature, from which we deduce the optical power of the amplified seeder beam and thus the amplification factor.

With seeding beam at 1 picowatt, we observe the amplification factor better than 20dB. With higher seeding power the amplification is reduced, but we still get a significant amount of amplification (better than 10dB) even at 100 picowatt level. Note that the amplification factor we use here is defined in terms of the optical power, not the field strength (i.e. quadrature values). The gain curve looks something like Lorentzian at 1 picowatt amplification, and the FWHM (full width at half maximum) is measured to be less than 0.5 MHz. At higher seeding powers, the FWHM becomes broader and the curve becomes highly asymmetric as is shown in the 100 picowatt

²In theory, such filtering is not necessary because we only extract information on the seeder beam frequency. However, because the pump beam signal is so strong it adds a significant amount of noise on the extracted amplitude/phase quadrature values.

case. Such distortion on the shape of the gain curve is common to all nonlinear resonance systems. (This is referred to as the “foldover effect”.)

7.3 Noise Analysis

From the gain curve measurement explained in the previous section, we found that the optimal seeder frequency that gave the maximum gain was -27.5 MHz detuned from the atomic resonance frequency. With the seeder frequency fixed at this optimal setting, we also have scanned the optical power of the seeder. The amplification factors measured as a function of the power of the seeding beam are marked with blue circles on the top plot of Figure 7.4. We already saw in the previous section that there is a trend of decreasing gain with increasing seeder power. Figure 7.4 demonstrates such correlation much more clearly; we see that the gain decreases in an almost linear fashion on the log-log scale over the range of input power that we tested.

It would have been great if we could continue to reduce the seeder power to verify that the low power gain maximizes to a constant (i.e. ~ 100). However, at tens of picowatt range the system receives less than one photon during one limit cycle period, and also the shot noise completely dominates the seeder beam. Therefore, the measurements become unreliable in those low power regimes.

Another interesting feature of our optical amplifier is its noise characteristics. It is well-known in the optical amplifier community that linear optical amplifiers suffer from the fundamental 3dB noise figure³, which is imposed by the quantum

³unless it is an optical degenerate parametric amplifier. In a strict sense, it is not an amplifier but a squeezer, since it amplifies only one of the quadratures and at the same time de-amplifies the other. Also note that degenerate parametric amplifiers are phase-sensitive, meaning that you need

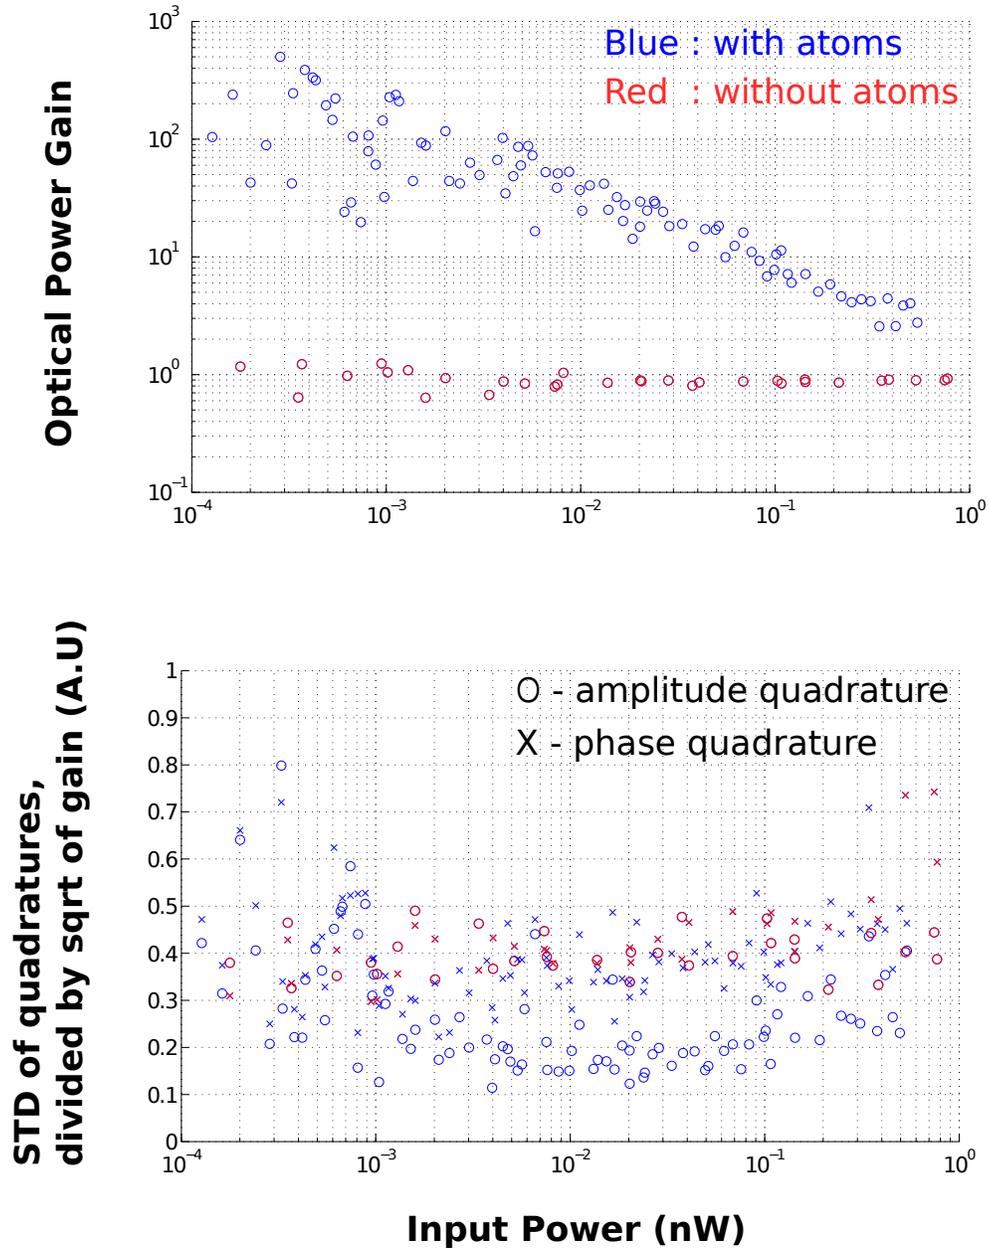


Figure 7.4: [TOP] Amplification factor as a function of input (seeder) power. Seeder frequency is chosen to be 27.5 MHz red-detuned from the atomic resonance frequency. Blue data is the result from the amplifier. Red data is a resonant, empty cavity response and therefore its amplification factor is a unity. [BOTTOM] Standard deviation on the measured quadrature values, divided by the square-root of the amplification factor.

mechanics [75]. In other words, when an ideal linear optical amplifier amplifies a coherent beam, say by 10dB, the noise on the amplified output is twice (3dB) as large as the initial noise level (shot noise, i.e. quantum-limited noise) amplified by the same amplification factor (10dB), creating 13dB amplification on the noise. This means that every time your optical signal is amplified by a linear amplifier, your signal to noise ratio (SNR) gets worse by a factor of two. The optical amplifier we demonstrate in this thesis, on the other hand, is a nonlinear amplifier and therefore its SNR characteristics does not need to be limited by such 3dB contamination.

In fact, the optical amplifier we demonstrated may have near-0dB noise figure, at least on the amplitude quadrature. This is hinted on the bottom plot of Figure 7.4, where we plotted the standard deviation of the measured quadrature values divided by the gain⁴, which is essentially the signal-to-noise ratio multiplied by some constant. Compared with unamplified coherent light (red), we see a very low noise level in the amplifier output (blue), most predominantly from 1 picowatt to 100 picowatt range.

Although the signal-to-noise ratio of the amplifier output in the figure seems to be below the coherent-beam level, it does not actually mean that the amplifier has a negative noise figure, which is logically impossible. There are two reasons why such error occurred when we produced this plot. The first reason is the optical loss on the cavity output, which causes the signal-to-noise ratio of the coherent beam to appear larger than it actually is. For a highly amplified beam, optical loss decreases the powers stored in both the signal and the noise by nearly the same factor, and therefore the signal-to-noise ratio of the beam stays relatively constant. But for a coherent beam, while the signal power gets reduced by the loss, the noise power does

to enforce the correct phase relation between the seeding beam and your system in order to get a good gain from your amplifier.

⁴Note that the gain in the quadrature space is the square root of the gain in optical power.

not. This is because the coherent beam already has the fundamentally lowest noise level, i.e. the quantum shot noise. It helps to understand this unbalanced effect of the optical loss on the signal-to-noise ratio if we interpret the optical loss as a mixing of the beam with a vacuum state by a beam splitter, as is commonly done in cavity QED theory. When the noise on your beam is much larger than the vacuum fluctuation, such mixing reduces the noise level, but if the original noise is already at the level of vacuum fluctuation (for example, if it is a coherent beam), mixing it with another vacuum fluctuation causes no reduction in the noise level.

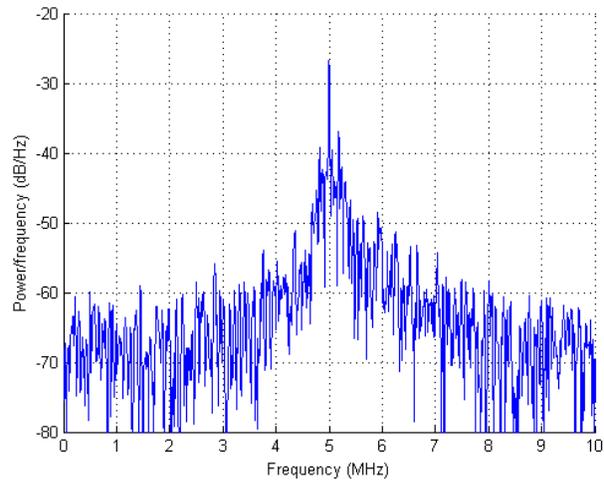
Another reason for the apparent negative noise figure⁵ is because we used, by mistake, the data traces that have too long temporal lengths while making the bottom plot of Figure 7.4. Each data point in this plot represents the standard deviation of the quadrature values that are extracted from 100 segments of 1 microsecond traces. We prepared these segments by fragmenting one continuous data trace from one experimental run which was originally 100 microsecond long. This is a perfectly fine method to use when we analyze the empty cavity output (coherent beam signal-to-noise ratio measurement, in red), because the our cavity linewidth is $2\pi \cdot 2.3$ MHz and therefore its decay time is much faster than 1 microsecond. However, the gain curve of the amplifier (Fig. 7.3) shows that it has less than 1 MHz bandwidth, meaning that the amplifier response is slower than 1 microsecond. We think that the standard deviations we plotted for the amplifier output (in blue) was probably under-estimated.

With these two points in mind, the exact value of the amplifier noise figure is still being investigated by Dr. Soh who is continuing the research and the experiment on this subject.

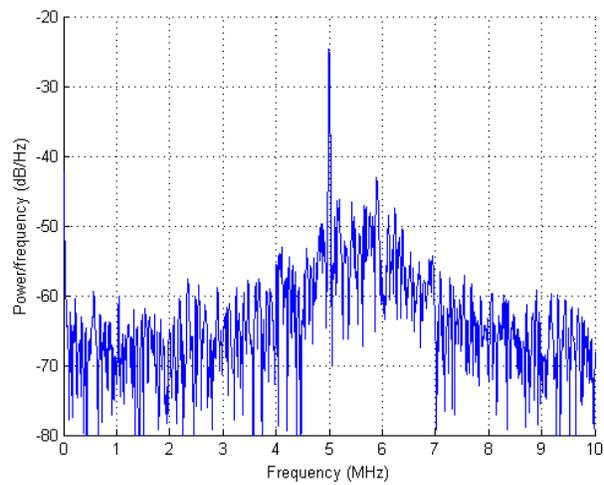
And finally, in Figure 7.5 we show the measured power spectrum of the amplifier

⁵This second reason was found while writing this thesis.

1 pW



10 pW



100 pW

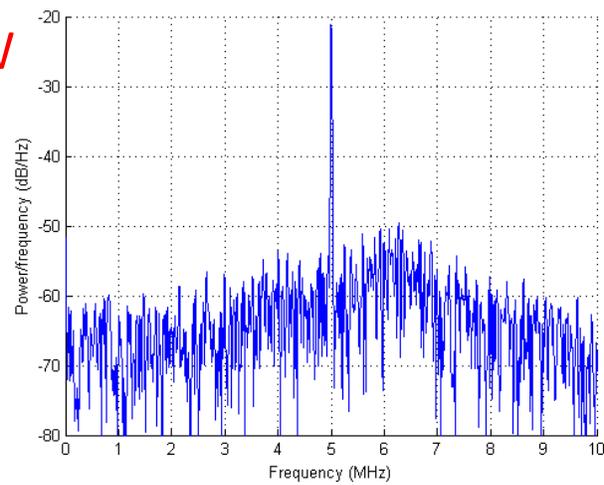


Figure 7.5: Decreasing amplifier output noise with larger seeding power

output. The local oscillator frequency was 5 MHz detuned from the the seeding beam, and therefore the sharp peak located at a position labeled as 5 MHz represents the measured amplified signal. The bump shown in the power spectrum around the peak translates to the level of noise contaminating the signal when we measure the beam quadratures. That is, the larger the bump size is, the higher the standard deviation the measured quadrature values will display. The three plots in the figure clearly show that the noise decreases with larger seeder power. This is not so surprising, as we already have demonstrated in Figure 7.4 that the noise figure of the amplifier is relatively constant at near 0 dB in this range of the seeder power. Because there is less amplification of the signal as we increase the seeder power, there should be less amplification of the noise (in this case, the amplified quantum shot noise) as well. However, these plots remind us of the behavior of a lasing gain medium, in which a strong drive at a particular frequency tends to suck out the stored energy from the whole gain spectrum, leaving less energy to be lased at other frequencies. The nonlinear optical amplifier we demonstrate here seems to have a similar behavior.

Bibliography

- [1] Paul Mandel. *Theoretical Problems in Cavity Nonlinear Optics*. Cambridge University Press, 2005.
- [2] Luigi A. Lugiato and E. Wolf. Theory of optical bistability. In *Progress in Optics*, volume 21, pages 69–216. Elsevier, 1984.
- [3] Kenju Otsuka. *Nonlinear Dynamics in Optical Complex Systems*. Springer, 2000.
- [4] Hyatt M. Gibbs. *Optical Bistability: Controlling Light with Light*. Academic Press, 1985.
- [5] G. Rempe, R. J. Thompson, R. J. Brecha, W. D. Lee, and H. J. Kimble. Optical bistability and photon statistics in cavity quantum electrodynamics. *Phys. Rev. Lett.*, 67(13):1727–1730, Sep 1991.
- [6] H. Mabuchi, Q. A. Turchette, M. S. Chapman, and H. J. Kimble. Real-time detection of individual atoms falling through a high-finesse optical cavity. *Opt. Lett.*, 21(17):1393–1395, Sep 1996.
- [7] J. A. Sauer, K. M. Fortier, M. S. Chang, C. D. Hamley, and M. S. Chapman. Cavity qed with optically transported atoms. *Phys. Rev. A*, 69:051804, May 2004.

- [8] Subhadeep Gupta, Kevin L. Moore, Kater W. Murch, and Dan M. Stamper-Kurn. Cavity nonlinear optics at low photon numbers from collective atomic motion. *Phys. Rev. Lett.*, 99:213601, Nov 2007.
- [9] Martin Mücke, Eden Figueroa, Joerg Bochmann, Carolin Hahn, Karim Murr, Stephan Ritter, Celso J. Villas-Boas, and Gerhard Rempe. Electromagnetically induced transparency with single atoms in a cavity. *Nature*, 465(7299):755–758, 06 2010.
- [10] Joseph Kerckhoff, Michael A. Armen, Dmitri S. Pavlichin, and Hideo Mabuchi. The dressed atom as binary phase modulator: towards attojoule/edge optical phase-shift keying. *Opt. Express*, 19(7):6478–6486, Mar 2011.
- [11] D. J. Alton, N. P. Stern, Takao Aoki, H. Lee, E. Ostby, K. J. Vahala, and H. J. Kimble. Strong interactions of single atoms and photons near a dielectric boundary. *Nat Phys*, 7(2):159–165, 02 2011.
- [12] Kartik Srinivasan and Oskar Painter. Linear and nonlinear optical spectroscopy of a strongly coupled microdisk-quantum dot system. *Nature*, 450(7171):862–865, 12 2007.
- [13] Andrei Faraon, Ilya Fushman, Dirk Englund, Nick Stoltz, Pierre Petroff, and Jelena Vuckovic. Coherent generation of non-classical light on a chip via photon-induced tunnelling and blockade. *Nat Phys*, 4(11):859–863, 11 2008.
- [14] Kengo Nozaki, Akihiko Shinya, Shinji Matsuo, Yasumasa Suzaki, Toru Segawa, Tomonari Sato, Yoshihiro Kawaguchi, Ryo Takahashi, and Masaya Notomi.

- Ultralow-power all-optical ram based on nanocavities. *Nat Photon*, 6(4):248–252, 04 2012.
- [15] Simon Groblacher, Klemens Hammerer, Michael R. Vanner, and Markus Aspelmeyer. Observation of strong coupling between a micromechanical resonator and an optical cavity field. *Nature*, 460(7256):724–727, 08 2009.
- [16] Stefan Weis, Rémi Rivière, Samuel Deléglise, Emanuel Gavartin, Olivier Arcizet, Albert Schliesser, and Tobias J. Kippenberg. Optomechanically induced transparency. *Science*, 330(6010):1520–1523, 2010.
- [17] Vivek Venkataraman, Kasturi Saha, and Alexander L. Gaeta. Phase modulation at the few-photon level for weak-nonlinearity-based quantum computing. *Nat Photon*, 7(2):138–141, 02 2013.
- [18] Howard J. Carmichael. *Statistical Methods in Quantum Optics 2: Non-Classical Fields*. Springer, 2007.
- [19] M. A. Armen and H. Mabuchi. Low-lying bifurcations in cavity quantum electrodynamics. *Phys. Rev. A*, 73:063801, Jun 2006.
- [20] Yurii A. Vlasov. Silicon cmos-integrated nano-photonics for computer and data communications beyond 100g. *IEEE Communications Magazine*, 50(2), Feb 2012.
- [21] C Santori, P E Barclay, K-M C Fu, R G Beausoleil, S Spillane, and M Fisch. Nanophotonics for quantum optics using nitrogen-vacancy centers in diamond. *Nanotechnology*, 21(27):274008, 2010.

- [22] David A. B. Miller. Are optical transistors the logical next step? *Nat Photon*, 4(1):3–5, 01 2010.
- [23] Yeong-Dae Kwon, Michael A. Armen, and Hideo Mabuchi. Femtojoule-scale all-optical latching and modulation via cavity nonlinear optics. *Phys. Rev. Lett.*, 111:203002, Nov 2013.
- [24] H. M. Gibbs, S. L. McCall, and T. N. C. Venkatesan. Differential gain and bistability using a sodium-filled fabry-perot interferometer. *Phys. Rev. Lett.*, 36:1135–1138, May 1976.
- [25] K. G. Weyer, H. Wiedenmann, M. Rateike, W. R. Mac Gillivray, P. Meystre, and H. Walther. Observation of absorptive optical bistability in a fabry-perot cavity containing multiple atomic beams. *Optics Communications*, 37(6):426–430, 6 1981.
- [26] A. T. Rosenberger, L. A. Orozco, H. J. Kimble, and P. D. Drummond. Absorptive optical bistability in two-state atoms. *Phys. Rev. A*, 43:6284–6302, Jun 1991.
- [27] P.D. Drummond. Optical bistability in a radially varying mode. *Quantum Electronics, IEEE Journal of*, 17(3):301–306, 1981.
- [28] R. J. Ballagh, J. Cooper, M. W. Hamilton, W. J. Sandle, and D. M. Warrington. Optical bistability in a gaussian cavity mode. *Optics Communications*, 37(2):143–148, 4 1981.
- [29] R. J. Thompson, G. Rempe, and H. J. Kimble. Observation of normal-mode splitting for an atom in an optical cavity. *Phys. Rev. Lett.*, 68:1132–1135, Feb 1992.

- [30] Yves Colombe, Tilo Steinmetz, Guilhem Dubois, Felix Linke, David Hunger, and Jakob Reichel. Strong atom-field coupling for bose-einstein condensates in an optical cavity on a chip. *Nature*, 450(7167):272–276, 11 2007.
- [31] Ferdinand Brennecke, Tobias Donner, Stephan Ritter, Thomas Bourdel, Michael Kohl, and Tilman Esslinger. Cavity qed with a bose-einstein condensate. *Nature*, 450(7167):268–271, 11 2007.
- [32] C. J. Hood, M. S. Chapman, T. W. Lynn, and H. J. Kimble. Real-time cavity qed with single atoms. *Phys. Rev. Lett.*, 80:4157–4160, May 1998.
- [33] Kengo Nozaki, Takasumi Tanabe, Akihiko Shinya, Shinji Matsuo, Tomonari Sato, Hideaki Taniyama, and Masaya Notomi. Sub-femtojoule all-optical switching using a photonic-crystal nanocavity. *Nat Photon*, 4(7):477–483, 07 2010.
- [34] H. M. Gibbs, F. A. Hopf, D. L. Kaplan, and R. L. Shoemaker. Observation of chaos in optical bistability. *Phys. Rev. Lett.*, 46:474–477, Feb 1981.
- [35] H. Nakatsuka, S. Asaka, H. Itoh, K. Ikeda, and M. Matsuoka. Observation of bifurcation to chaos in an all-optical bistable system. *Phys. Rev. Lett.*, 50:109–112, Jan 1983.
- [36] F. Mitschke, J. Mlynek, and W. Lange. Observation of magnetically induced optical self-pulsing in a fabry-perot resonator. *Phys. Rev. Lett.*, 50:1660–1663, May 1983.
- [37] Bernard Segard and Bruno Macke. Self-pulsing in intrinsic optical bistability with two-level molecules. *Phys. Rev. Lett.*, 60:412–415, Feb 1988.

- [38] L. A. Lugiato, R. J. Horowicz, G. Strini, and L. M. Narducci. Instabilities in passive and active optical systems with a gaussian transverse intensity profile. *Phys. Rev. A*, 30:1366–1376, Sep 1984.
- [39] L. A. Orozco, A. T. Rosenberger, and H. J. Kimble. Intrinsic dynamical instability in optical bistability with two-level atoms. *Phys. Rev. Lett.*, 53:2547–2550, Dec 1984.
- [40] L. A. Orozco, H. J. Kimble, A. T. Rosenberger, L. A. Lugiato, M. L. Asquini, M. Brambilla, and L. M. Narducci. Single-mode instability in optical bistability. *Phys. Rev. A*, 39:1235–1252, Feb 1989.
- [41] Victor Grigoriev and Fabio Biancalana. Resonant self-pulsations in coupled nonlinear microcavities. *Phys. Rev. A*, 83:043816, Apr 2011.
- [42] Stefania Malaguti, Gaetano Bellanca, Alfredo de Rossi, Sylvain Combrié, and Stefano Trillo. Self-pulsing driven by two-photon absorption in semiconductor nanocavities. *Phys. Rev. A*, 83:051802, May 2011.
- [43] T. Gu, N. Petrone, J. F. McMillan, A. van der Zande, M. Yuand G. Q. Lo, D. L. Kwong, J. Hone, and C. W. Wong. Regenerative oscillation and four-wave mixing in graphene optoelectronics. *Nat Photon*, 6(8):554–559, 08 2012.
- [44] J. D. Weinstein and K. G. Libbrecht. Microscopic magnetic traps for neutral atoms. *Phys. Rev. A*, 52(5):4004–4009, Nov 1995.
- [45] J. Reichel, W. Hänsel, and T. W. Hänsch. Atomic micromanipulation with magnetic surface traps. *Phys. Rev. Lett.*, 83:3398–3401, Oct 1999.

- [46] J. Reichel, W. Hänsel, P. Hommelhoff, and T. W. Hänsch. Applications of integrated magnetic microtraps. *Applied Physics B: Lasers and Optics*, 72(1):81–89, 01 2001.
- [47] Ron Folman, Peter Krüger, Jörg Schmiedmayer, Johannes Denschlag, Carsten Henkel, Benjamin Bederson, and Herbert Walther. *Microscopic atom optics: From wires to an atom chip*, volume Volume 48, pages 263–356. Academic Press, 2002.
- [48] Benjamin Lev, Kartik Srinivasan, Paul Barclay, Oskar Painter, and Hideo Mabuchi. Feasibility of detecting single atoms using photonic bandgap cavities. *Nanotechnology*, 15(10):S556–S561, 2004.
- [49] Benjamin L. Lev. *Magnetic Microtraps for Cavity QED, Bose-Einstein condensates, and Atom Optics*. California Institute of Technology, <http://resolver.caltech.edu/CaltechETD:etd-09202005-205733>, 2005.
- [50] Pascal Bohi, Max F. Riedel, Johannes Hoffrogge, Jakob Reichel, Theodor W. Hänsch, and Philipp Treutlein. Coherent manipulation of bose-einstein condensates with state-dependent microwave potentials on an atom chip. *Nat Phys*, 5(8):592–597, 08 2009/08//print.
- [51] J. Reichel, W. Hänsel, P. Hommelhoff, and T.W. Hänsch. Applications of integrated magnetic microtraps. *Applied Physics B*, 72(1):81–89, 2001.
- [52] Ron Folman, Peter Krüger, Donatella Cassettari, Björn Hessmo, Thomas Maier, and Jörg Schmiedmayer. Controlling cold atoms using nanofabricated surfaces: Atom chips. *Phys. Rev. Lett.*, 84:4749–4752, May 2000.

- [53] M. P. A. Jones, C. J. Vale, D. Sahagun, B. V. Hall, and E. A. Hinds. Spin coupling between cold atoms and the thermal fluctuations of a metal surface. *Phys. Rev. Lett.*, 91(8):080401, Aug 2003.
- [54] M.-O. Mewes, M. R. Andrews, D. M. Kurn, D. S. Durfee, C. G. Townsend, and W. Ketterle. Output coupler for bose-einstein condensed atoms. *Phys. Rev. Lett.*, 78:582–585, Jan 1997.
- [55] B. P. Anderson and M. A. Kasevich. Macroscopic quantum interference from atomic tunnel arrays. *Science*, 282(5394):1686–1689, 1998.
- [56] Immanuel Bloch, Theodor W. Hänsch, and Tilman Esslinger. Atom laser with a cw output coupler. *Phys. Rev. Lett.*, 82:3008–3011, Apr 1999.
- [57] W. Hänsel, J. Reichel, P. Hommelhoff, and T. W. Hänsch. Magnetic conveyor belt for transporting and merging trapped atom clouds. *Phys. Rev. Lett.*, 86:608–611, Jan 2001.
- [58] E. Torrontegui. Fast atomic transport without vibrational heating. *Phys. Rev. A*, 83(013415), 2011.
- [59] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping (Graduate Texts in Contemporary Physics)*. Springer, November 2001.
- [60] J. Reichel. Microchip traps and bose-einstein condensation. *Applied Physics B: Lasers and Optics*, 74(6):469–487, 04 2002/04/15/.
- [61] R. Long, T. Rom, W. Hänsel, T. W. Hänsch, and J. Reichel. Long distance magnetic conveyor for precise positioning of ultracold atoms. *The European*

- Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 35(1):125–133, 2005.
- [62] S. L. McCall. Instabilities in continuous-wave light propagation in absorbing media. *Phys. Rev. A*, 9:1515–1523, Apr 1974.
- [63] K. Hennessy, A. Badolato, M. Winger, D. Gerace, M. Atature, S. Gulde, S. Falt, E. L. Hu, and A. Imamoglu. Quantum nature of a strongly coupled single quantum dot-cavity system. *Nature*, 445(7130):896–899, 02 2007.
- [64] Paul E. Barclay, Kai-Mei Fu, Charles Santori, and Raymond G. Beausoleil. Hybrid photonic crystal cavity and waveguide for coupling to diamond nv-centers. *Opt. Express*, 17(12):9588–9601, Jun 2009.
- [65] Ryan M. Briggs, Michael Shearn, Axel Scherer, and Harry A. Atwater. Wafer-bonded single-crystal silicon slot waveguides and ring resonators. *Applied Physics Letters*, 94(2):021106, 2009.
- [66] Hansuek Lee, Tong Chen, Jiang Li, Ki Youl Yang, Seokmin Jeon, Oskar Painter, and Kerry J. Vahala. Chemically etched ultrahigh-q wedge-resonator on a silicon chip. *Nat Photon*, 6(6):369–373, 06 2012.
- [67] Dirk Englund, Arka Majumdar, Andrei Faraon, Mitsuru Toishi, Nick Stoltz, Pierre Petroff, and Jelena Vučković. Resonant excitation of a quantum dot strongly coupled to a photonic crystal nanocavity. *Phys. Rev. Lett.*, 104:073904, Feb 2010.
- [68] H. John Caulfield and Shlomi Dolev. Why future supercomputing requires optics. *Nat Photon*, 4(5):261–263, 05 2010.

- [69] Wenge Yang, Amitabh Joshi, and Min Xiao. Chaos in an electromagnetically induced transparent medium inside an optical cavity. *Phys. Rev. Lett.*, 95:093902, Aug 2005.
- [70] Bjorn Maes, Martin Fiers, and Peter Bienstman. Self-pulsing and chaos in short chains of coupled nonlinear microcavities. *Phys. Rev. A*, 80:033805, Sep 2009.
- [71] Steven H. Strogatz. *Nonlinear Dynamics And Chaos: With Applications To Physics, Biology, Chemistry, And Engineering*. Studies in nonlinearity. Perseus Books Group, 1 edition, January 1994.
- [72] Paul E. Barclay, Kai-Mei C. Fu, Charles Santori, and Raymond G. Beausoleil. Chip-based microcavities coupled to nitrogen-vacancy centers in single crystal diamond. *Applied Physics Letters*, 95(19):191115, 2009.
- [73] Sunil Sandhu, Michelle L. Povinelli, and Shanhui Fan. Enhancing optical switching with coherent control. *Applied Physics Letters*, 96(23):231108, 2010.
- [74] Kurt Wiesenfeld and Bruce McNamara. Small-signal amplification in bifurcating dynamical systems. *Phys. Rev. A*, 33:629–642, Jan 1986.
- [75] Carlton M. Caves. Quantum limits on noise in linear amplifiers. *Phys. Rev. D*, 26:1817–1839, Oct 1982.