Continuous observation of quantum dynamics

Thesis by

Hideo Mabuchi

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

1998

(Submitted April 29, 1998)

Acknowledgements

I believe there is only one way that a student can ever repay the invaluable gift of guidance and inspiration from a great teacher—the student must endeavor to become a great teacher as well. So my true thanks to Jeff Kimble will be many years in the making. Nonetheless, let me here record my sincere gratitude to him for teaching me experimental and theoretical physics, for fostering my development as a creative researcher, and for impressing upon me the grave importance of valuing scientific integrity above all other standards for professional self-assessment.

I would like to acknowledge the crucial importance of my early interactions with Artur Ekert and Charlie Bennett. Their earnest confidence and interest in my thoughts enabled me to define myself as an independent scientist long before I could have done so on my own. I must also thank Eugene Polzik, Vladimir Ilchenko, and Peter Zoller for teaching me through collaboration and by example, both in the lab and on the napkin. In a less direct but equally important way, I have also been shaped and guided by Vladimir Braginsky and Carlton Caves. They have had a profound influence on the nature of my ambitions in science. Working in the lab with Quentin Turchette was one of the highlights of my time as a graduate student, and my recent experience of working with Jun Ye has been another. Many thanks to them both, and many thanks as well to Christopher Fuchs for our endless hours of musing conversation on truth, chance, and the quantum.

Finally, I would like to thank some colleagues that I have had the good fortune to work with over the years at Caltech: Nikos Georgiades, Rob Thompson, Keiichi Edamatsu, Oliver Carnal, John Hartman, Yuri Levin, Wolfgang Lange, Christina Hood, David Vernooy, Mike Chapman, Theresa Lynn, Akira Furusawa, Joe Buck, Steven Van Enk, Ken Libbrecht, and John Preskill.

Abstract

Previous work in cavity quantum electrodynamics has focused principally on achieving the condition of strong coupling, such that the single-photon Rabi frequency for one atom in a cavity dominates the rates associated with all forms of dissipation and dephasing. Although strong coupling was convincingly demonstrated in these experiments, the use of thermal atomic beams as a source of atoms invariably limited the intracavity dwell time of any given atom to values marginally longer than the system's characteristic dynamical timescales. Accordingly, the observation of single-atom effects in these experiments required integration over the transits of many successive atoms through the optical resonator. In such ensemble-averaged experiments, the quantitative effects of dissipation and measurement could always be understood from the perspective of Master Equation formalisms.

By contrast, our most recent work has incorporated laser cooling techniques to accomplish a thousand-fold increase in the dwell time of individual atoms while maintaining conditions of strong coupling. With this extreme separation of timescales, the effects of individual atoms on the cavity transmission can clearly be resolved. Moreover, detailed information on each individual atom's spatial trajectory through the cavity can be inferred from the time-varying amplitude and phase of the transmitted light. Understanding the exact nature of such correlations in real laboratory measurements requires not only a careful characterization of excess technical noise, but also an in-depth study of how (and when) aspects of continuous quantum measurement theory should properly be applied.

This thesis primarily contains a detailed description of an experiment in which we used broadband heterodyne spectroscopy to record the complete time-evolution of interaction energy between one atom and a high-finesse optical cavity, during individual scattering events of ~ 250 μs duration. With a characteristic interaction energy scale $E_{\rm int}/\hbar \sim 10$ MHz, we achieve a shot-noise limited measurement sensi-

tivity $\simeq 3$ kHz/$\sqrt{\text{Hz}}$ over a bandwidth that covers the dominant rates of variation in E_{int}. I also include two brief studies in quantum measurement theory, which highlight the significance of the experimental work for continuing research on conditional quantum dynamics. The first investigates quantum limits to broadband measurement of position and force, and the second proposes a novel trajectory-based approach to quantum system identification.

Contents

List of Figures

Preface

"Six years in Kimble-group"

Upon arriving at Caltech during the summer of 1992, I was kindly greeted by my new advisor but then summarily instructed to "make the rounds" and find myself a research project. By some good fortune I ended up working in the lab with Eugene Polzik, helping him and Nikos Georgiades in the very early stages of their experiment on two-photon spectroscopy with squeezed light. I have Eugene to thank for quickly bringing me up to speed with the basic methods of experimental quantum optics. My year-long stage in the squeezing lab culminated in a set of measurements on blue-light induced infrared absorption in potassium niobate, which we wrote up as a paper in the Journal of the Optical Society of America [1].

Although it seemed clear at the time that the two-photon squeezing experiment was building up momentum and would be destined for scientific glory, I went to Jeff Kimble at the end of that first year with a request to forge out on an experiment of my own. Much to my amazement, he took this sophomoric uprising in stride and proposed that I undertake the development of quartz microsphere optical resonators as a new technical paradigm for experiments in cavity quantum electrodynamics. That sounded great.

My first task in this new research program was to piece together some understanding of the optical mode structure of microspheres, and then to compute the set of parameters that would tell us whether microspheres could really be useful for cavity QED. After several months I had numbers in hand, and they were sufficiently promising that we decided to pack our bags and fly to Moscow for a hands-on lesson in microsphere fabrication techniques with the expert members of Vladimir Braginsky's precision measurement group. After two weeks of practice in Russia with Vladimir Ilchenko and Michael Gorodetsky, I returned to build up a microsphere effort of our

own at Caltech.

The following year brought a number of important developments, beginning with my first success in reproducing the record-high optical quality factors ($\sim 2 \times 10^9$) demonstrated in Moscow. The push towards even higher quality factors has been carried on by David Vernooy for the past few years, and he was kind enough to acknowledge my early contributions to the effort by including me as a co-author on his recent Optics Letter [2]. That spring, Vladimir Ilchenko came to visit Caltech for a collaborative attempt to see some actual cavity-QED signals using quartz microspheres. We didn't manage to produce any data during Vladimir's visit, but we did accomplish enough to prove that microspheres were indeed a viable technology for experiments on strong coupling with atoms. In the post-mortem analysis of our first failed attempt at doing so, I became convinced that one needed to have a *very* cold and dense sample of atoms in order to really do cavity QED with microspheres. Although Dave has since proven me wrong, I'm glad that my beliefs at the time led me to begin working on an apparatus in which I could form a magnetooptic trap of Cesium atoms in close proximity to a microsphere and all its associated optics and actuators.

In parallel with these experimental activities, Jeff and I put some serious effort into theoretical studies of what unique experiments could be enabled by microsphere technology. Our two most interesting proposals were written up as an Optics Letter [3], and an article in Physics Letters A that we wrote together with Russian colleagues Sergei Vyatchanin and Andrei Matsko [4]. The Optics Letter describes a scheme for confining cold atoms in orbit around a microsphere, using two optical fields of disparate wavelength to establish a stable mechanical potential for atomic center-of-mass motion. The Physics Letter presents a calculation demonstrating the feasibility of quantum non-demolition detection of single photons in a microsphere via quantized deflection of an atomic beam.

Just after I had seen my first Cesium trap glowing brightly in the center of the would-be microsphere vacuum chamber (some time in the summer of 1994), Scott Parkins came to visit our group and told us about some calculations he had done on

all the great physics one could do by dropping cold atoms into a Fabry-Perot optical cavity. After hearing Scott's results, it seemed that we could be poised for a quick kill by setting aside the microsphere business for a few months and putting a Fabry-Perot under the trap instead. With this motivation, Quentin Turchette and I started working together on this "little detour."

Also during the summer of 1994 we had our first visit from Artur Ekert, who told us about Peter Shor's very exciting but as-yet unverified quantum factoring algorithm. Over the course of dinner, Jeff and Artur and I came to realize that cavity QED had all the right qualities as a candidate for experimental investigations of quantum computing. In addition to working in the lab with Quentin, my autumn and winter of 1994-95 were thus largely occupied by working with Jeff to properly develop the exact relationship between strong coupling in cavity QED and *bona fide* quantum logic. By the following summer Quentin and Christina Hood had made some beautiful measurements, and together with Wolfgang Lange (who had been helping out on all fronts) we wrote up the Physical Review Letter that officially launched our group into the quantum computing fray [5].

In between all the quantum logic business, Quentin and I kept plugging away at our cavity QED with cold atoms project. Gradually we came to realize that the experiment was much more involved than we had originally appreciated, and around the beginning of summer 1995 we found ourselves ready to start over and redesign things from scratch. So while we waited for new mirrors and a new cavity mount to be fabricated and delivered, Quentin worked on other things and I went to spend six weeks at the Ecole de Physique d'Ete des Houches. While there, I had a conversation with Adriano Barenco and Peter Zoller that eventually led to a theory paper on inversion of quantum jumps, which appeared in Physical Review Letters in early 1996 [6]. With the ideas in that paper we thought we had made some real progress towards quantum error correction, although we were to be upstaged in that regard by Peter Shor's paper that appeared while ours was still in press.

During the fall of 1995, Mike Chapman joined Quentin and me in our laboratory efforts. Mike helped us out in developing our *third*-generation apparatus for cavity

QED with cold atoms. Late one night in March of 1996, we finally saw our first real-time atom detection signals and immediately faxed them off to Jeff, who was at a quantum measurement workshop in Germany. Luckily we caught him before his talk, so the world got to learn of our results just as we were going to bed. Thinking (for the second time) that a steady stream of further results must surely be just around the corner, we decided to take what we had and quickly wrote up an Optics Letter to put our data in print [7].

In fact, it seemed that there was so much good work to be done with our newly-proven system that we as a group decided to build a second experimental apparatus of the same general ilk and pursue two different objectives in parallel. So Mike went to work with Christina on a new lab, while Quentin and I set ourselves to the task of following up on some preliminary indications in our data that we might have been resolving atomic center-of-mass motion with precision near the Standard Quantum Limit. As evidenced by certain sections of this thesis, Quentin and I had no idea at the time of just what we were getting ourselves into.

In addition to doing laboratory work on characterizing noise sources and improving all the servos, I spent some time during early 1996 thinking about the quantum theory of atomic position measurements via cavity QED. Having already been exposed to quantum trajectory theory through my work with Peter Zoller, I adopted this as the natural mathematical tool for connecting what I knew from physics about the optical response of a strongly-coupled atom and cavity to what I was learning in statistics and information theory about optimal signal recovery. Some of my early results in exploring this connection are described in a Letter to the Editor in Quantum and Semiclassical Optics [8]. I rejoined the subject during a brief stay at the Institute for Theoretical Physics in Santa Barbara during October of 1996, focusing principally on the question of what the Standard Quantum Limit really meant in the context of *broadband* measurements such as those we were performing in the lab. Also during the ITP, Jeff and I had some extended discussions with Peter Zoller and Ignacio Cirac that eventually led to the publication of a theoretical proposal for quantum state transfer and entanglement distribution between atoms located in spatially-distant cavities [9].

In terms of experimental work, we spent the second half of 1996 making endless incremental improvements to the apparatus. The main objective was to monitor atomic motion with a detuned probe beam, in order to circumvent the heating mechanisms that Andrew Doherty and our other colleagues at the University of Auckland helped us to realize were a significant problem. This entailed a complete revamping of the way we processed the rf photocurrent, and also demanded that we somehow achieve a drastic increase in our signal-to-noise ratio. Despite our valiant efforts, the progress was still hard to measure by the time that January 1997 rolled around and Quentin had to finish up and write his thesis.

At the time I was actually eager to work for a while on my own again, which I had not had the chance to do since the microsphere era. It was clear to me that the laser-frequency and cavity-length stabilization servos had to get much more serious than was typical for the group in those days, and that the only way for this to happen was for me to spend five or six months learning about feedback control the hard way. So after completely rebuilding the experiment, I systematically went through every feedback loop, measuring the open-loop transfer functions and designing servo amplifiers for optimal performance. Gradually, some real improvements in the signal-to-noise ratio started to add up, and I was finally able to produce some preliminary versions of the kind of data that will be presented in subsequent chapters.

During the summer of 1997 I travelled to several workshops on quantum computation in Italy. While there, Christopher Fuchs and I worked together and had numerous discussions with other attendees on the subject of how to quantify the entanglement of bipartite mixed states. As an outgrowth of this initial work, we now have a many-author paper in preparation on one particular measure that we investigated, called the Entanglement of Assistance [10].

After returning from Italy I was joined in the lab by our new postdoctoral fellow, Jun Ye. Thanks largely to Jun's formidable expertise in optical techniques for precision measurement, the experiment really came together during the final months of 1997. Through a sequence of final and decisive improvements to the apparatus, we were able finally to push the signal-to-noise ratio for monitoring single-atom motion

all the way down to the quantum limit. We took most of the data that I will discuss in this thesis during December of 1997, although we continued to work on a number of important calibrations and also attempted some further improvements to the experiment through February of 1998. In conjunction with the writing of this thesis, I am preparing an extended article on our apparatus and experimental results, which will most likely be submitted to Applied Physics B [11].

As a result of some very stimulating discussions with Salman Habib earlier in 1997, I made two visits to the Los Alamos National Laboratory during September and November. In order to lay the groundwork for our ongoing collaborations on quantum chaos and quantum feedback control, Salman introduced me to the pertinent aspects of high-performance computing and enabled me to start writing codes and running jobs on some very powerful machines. I am indebted to him for making possible most of the numerical results that are included in this thesis. I should also mention Sze Tan's visit to Caltech in the early months of 1996 as having made a profound contribution to my capabilities for computational research.

Over the 1997 Christmas break, and during a workshop in early January 1998, I finally finished up my theoretical work on characterizing quantum limits to broadband position measurements. I have recently written them up as a paper that will appear in Physical Review A [12].

In choosing the material for inclusion in this thesis, I have decided to focus on the most recent work that I have done as a graduate student at Caltech. As indicated in the listing below, my earlier research has been reasonably well documented in refereed publications. Since our 1996 Optics Letter, my experimental and theoretical research have really grown together to form the basis of a single long-term scientific program. So with a view towards the future, and without any apology for the brevity of the resulting document, I have collected together only the work in which I take the most pride and present them under the title of *Continuous observation of quantum dynamics*. Having made this choice of emphasis, descriptions of my research on microspheres and on quantum computing are notably omitted. I hope the impression will not arise that I found my work in these areas any less formative or less gratifying,

but I do feel that I personally identify much more strongly with my efforts to create a foundation for investigating the dynamical evolution of continuously-observed open quantum systems. It is my strong conviction that this constitutes a singularly important scientific program for future research in cavity QED, so I hope that I will have much more to say about it in years to come.

Publications based on graduate work:

1. H. Mabuchi, E. S. Polzik, and H. J. Kimble, "Blue-light induced infrared absorption in $KNbO_3$," J. Opt. Soc. Am. B **11**, 2023-2029 (1994).

2. D. W. Vernooy, V. S. Ilchenko, H. Mabuchi, E. W. Streed, and H. J. Kimble, "High-Q measurements of fused-silica microspheres in the near-infrared," Opt. Lett. **23**, 247-249 (1998).

3. H. Mabuchi and H. J. Kimble, "Atom galleries for whispering atoms: binding atoms in stable orbits around a high-finesse optical resonator," Opt. Lett. **19**, 749-751 (1994).

4. A. B. Matsko, S. P. Vyatchanin, H. Mabuchi, and H. J. Kimble, "Quantum nondemolition detection of individual photons by atomic beam deflection," Phys. Lett. A **192**, 175-179 (1994).

5. Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble, "Measurement of conditional phase shifts for quantum logic," Phys. Rev. Lett. **75**, 4710-4713 (1995).

6. H. Mabuchi and P. Zoller, "Inversion of quantum jumps in open quantum-optical systems under continuous observation," Phys. Rev. Lett. **76**, 3801-3804 (1996).

7. 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**, 1393-1395 (1996).

*__8.__ H. Mabuchi, "Dynamical identification of open quantum systems," Quantum Semiclass. Opt. **8**, 1103-1108 (1996).

9. J.-I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi, "Quantum-state transfer and entanglement distribution among distant nodes in a quantum network," Phys. Rev. Lett. **78**, 3221-3224 (1997).

*__10.__ H. Mabuchi, J. Ye, and H. J. Kimble, "Full observation of single-atom dynamics in cavity QED," in preparation.

11. D. P. DiVincenzo, C. A. Fuchs, H. Mabuchi, M. A. Nielsen, J. A. Smolin, A. Thapliyal, and A. Uhlmann, "Entanglement of Assistance," in preparation.

*__12.__ H. Mabuchi, "Standard quantum limits for broadband position measurement," to appear in Physical Review A (1998).

* *Indicates a publication on which this thesis is substantially based.*

Chapter 1 Introduction and overview

1.1 Motivation: conditional quantum dynamics

What happens *after* quantum measurement?

If your physics education was anything like mine, the answer should be obvious and read something like this: "Having measured a Hermitian observable $\hat{O}$ and obtained the outcome (eigenvalue) λ_i, the system is projected onto the eigenstate $\vec{\psi}_i$ such that

$$\hat{O}\vec{\psi}_i = \lambda_i \vec{\psi}_i. \tag{1.1}$$

Subtle things happen if λ_i is degenerate, but nobody really worries about that."

Those who have thought more deeply on the subject will recognize that such an answer completely misses the point, in that it overlooks the very important and complexifying issue of partial measurements made on composite quantum systems. Suppose for example that we have an atom coupled to electromagnetic fields, and that we wish to make a "measurement" of its internal state by placing a photon counter to intercept some small fraction of the outgoing solid angle for spontaneous emission. The atom should undergo radiative decay if and only if it is in an excited state, but we can only see the photon if it is emitted in the right direction. What *exactly* are we allowed to say about the *quantum* state of the atom at the end of an interval of time in which we *don't* see a photon? Escalated to only slightly more complex scenarios, this sort of issue lies at the heart of many pressing questions in the fields of quantum measurement, decoherence, and quantum-classical correspondence.

It has long been known that the statement of equation (1.1) can be generalized in the following manner [1, 2]. Suppose that the quantum system of interest lives in a Hilbert space $\mathcal{H}_A$, and is coupled (either by design or by nature) to some ancillary system that lives in a Hilbert space $\mathcal{H}_B$. Let the overall Hamiltonian for the system,

ancilla, and the coupling between them be $\mathcal{H}_{\rm tot}$. If we can somehow assure that the initial joint state factors as $\rho_0^{\rm tot} = \rho_0^A \otimes \rho_0^B$, then the effect on the *system* state of evolving $\rho^{\rm tot}$ under $\mathcal{H}_{\rm tot}$ for some time and then performing a projective measurement on the *ancilla* can always be written as

$$\rho_0^A \to \hat{c}_i \rho_0^A \hat{c}_i^\dagger. \tag{1.2}$$

Here i is an index associated with the outcome of the measurement made on $\mathcal{H}_B$, and the $\hat{c}_i$ are a set of operators on $\mathcal{H}_A$. A moment's reflection reveals that the operators $\hat{c}_i$ should simply be given by

$$\hat{c}_i = \hat{\Pi}_i^B \exp\left(-i\mathcal{H}_{\rm tot}\tau/\hbar\right), \tag{1.3}$$

where τ is the measurement interval and $\hat{\Pi}_i^B$ is a projector onto the $i^{\rm th}$ eigenstate of whatever observable we measure on $\mathcal{H}_B$. The probability of the $i^{\rm th}$ outcome is thus given by

$$\Pr_i = {\rm Tr}\left[\rho_0^A \hat{c}_i^\dagger \hat{c}_i\right]. \tag{1.4}$$

The generalized "collapse" represented by equation (1.2) is commonly known as a *quantum operation* [3], and each operator $\hat{c}_i$ is called an *operation element.* It is worth noting that a set of collapse operators $\{\hat{c}_i\}$ that act as in (1.2) can still be derived if we choose to make generalized (as opposed to projective) measurements on the ancilla. In such cases the number of operation elements will be larger than the dimension of $\mathcal{H}_B$, but one always has $\sum_i \hat{c}_i^\dagger \hat{c}_i = \hat{1}_A$ (the identity operator on $\mathcal{H}_A$).

If for whatever reason it is not possible to distinguish among some subset $\mathcal{S}$ of the possible measurement outcomes i, we must average incoherently over the associated operations in order to derive the evolution of the system state:

$$\rho_0^A \to \sum_{i\in\mathcal{S}} \hat{c}_i \rho_0^A \hat{c}_i^\dagger. \tag{1.5}$$

In the extreme, $\mathcal{S}$ could correspond to the entire range of measurement outcomes—

this would be the case if instead of making a measurement on the ancilla we were just to throw it away after jointly evolving under $\hat{H}_{\rm tot}$. The choice of a particular basis for the $\hat{c}_i$ would then be immaterial, and the nonselective evolution (1.5) would correspond to tracing over $\mathcal{H}_B$ [4].

Operationally, the ancilla in the above discussion could be an auxilliary quantum "device" that we introduce for the purpose of making an indirect measurement on the system. This type of scheme is generally required for making quantum non-demolition measurements [5], and also for implementing optimal measurements in the contexts of quantum communication and quantum cryptography [6]. In an indirect measurement we gain information about the system to the extent that the probabilities $\Pr_i$ depend on ρ_0^A, and (1.2) or (1.5) represents the *measurement backaction*. Note that if we know the measurement outcome i precisely, we also have the fullest possible knowledge of the evolution of the system's quantum state.

In contrast to the direct measurement scenario of (1.1), however, the state of the system following an indirect measurement generally depends both on the outcome i and the pre-measurement state ρ_0^A. Hence the evolution described by (1.2) is not so much a *collapse* of the system state as a *disturbance* of it. As investigated in references [10, 11], it should even be possible to implement "measurement" $\hat{c}_i$'s that are nontrivial but unitary, and therefore reversible. Not surprisingly, it turns out that quantum operations of this type don't actually yield any information about the system, in that the $\Pr_i$ are completely independent of ρ_0^A. In general, the information gained by performing a quantum operation necessarily comes at the cost of increasing severity of the disturbance. For low-dimensional systems, Fuchs and Peres [12] have even derived inference-disturbance relations similar to the Heisenberg uncertainty principle. Hence, one can appreciate that the traditional binary opposition between destructive and quantum non-demolition measurements represents only two isolated points in a continuous spectrum of possible inference-disturbance tradeoffs. A primary challenge for measurement science in the coming years will thus be to elaborate the properties and utility of the most general quantum operations, and concurrently to build the experimental capabilities required to implement them in the lab.

Beyond quantum measurement, evolutions of the form (1.2,1.5) also play a central role in the dynamics of open quantum systems. With reference once again to the basic quantum operation scenario, note that the ancilla $\mathcal{H}_B$ could represent an environmental reservoir from which we are unable fully to isolate the system of interest. The traditional picture of decoherence in Markovian open quantum systems can in fact be recovered [4, 7] by tracing over the reservoir and resetting ρ_0^B to some fiducial (*e.g.*, vacuum or thermal) state once per reservoir correlation time. One should take particular care in dealing with an infinite-dimensional ancilla such as the electromagnetic field [8, 9], but quantum-optical Master Equations may fairly be regarded as a quasi-continuous limit of (1.5).

Should there not also be a differential equation associated to the unaveraged quantum operation (1.2)? Loosely speaking, the state of an open quantum system decoheres because available information about its evolution is not experimentally *recovered* from the reservoir. The growth of (von Neumann) entropy implied by the form of most quantum-optical Master Equations stems from the fact that the terms in the sum of (1.5) will generally have small inner products with one another. Given an open quantum system in the laboratory, it should be possible to reduce our uncertainty about the evolving system state if we can manage to perform measurements on the reservoir with sufficiently high bandwidth. We believe from recent work on quantum trajectory theory [8, 7, 13] that the dynamical evolution of such a continuously-observed open quantum system should be described by a Stochastic Schrödinger Equation (SSE). A typical form for such an equation would be

$$d\vec{\psi}_c(t) = \left\{ \frac{-i}{\hbar}\hat{H}_0 dt - \frac{1}{2}\sum_{j=1}^{N_c}\left(\hat{c}_j^\dagger \hat{c}_j - \langle \hat{c}_j^\dagger \hat{c}_j \rangle_c\right) dt + \sum_{j=1}^{N_c}\left(\frac{\hat{c}_j}{\sqrt{\langle \hat{c}_j^\dagger \hat{c}_j \rangle_c}} - 1\right) dN_j(t) \right\} \vec{\psi}_c(t). \tag{1.6}$$

Here N_c is the total number of reservoirs coupled to the system, $\hat{c}_j$ is a "jump operator" for the j^{th} reservoir (analogous to the operation elements introduced above), and $dN_j(t)$ is a stochastic increment. In qualitative terms, the numerical value of $dN_j(t)$ (here restricted to be either 0 or 1) during any given timestep dt corresponds

to the outcome of a measurement made on the j^{th} reservoir at time t. Hence, $\vec{\psi}_c$ represents an *a posteriori* quantum state for the open system, conditioned upon the classical record that we obtain from our measurements on the reservoir. At the same time, one should keep in mind that the statistics of these measurements are determined by the expectation values of the operators $\hat{c}_j^\dagger \hat{c}_j$ with respect to $\vec{\psi}_c$. This type of dynamics for an open quantum system, in which the evolution of the system state is conditioned upon the measurement results *and vice versa*, is what I mean by the phrase "conditional quantum dynamics."

Starting from the Stochastic Schrödinger Equation for a highly idealized, but not completely unrealistic model of certain experimental systems that have recently been investigated in quantum optics, one can formulate all kinds of interesting schemes that would utilize continuous observation to accomplish novel quantum measurements [14, 16], to build feedback-control devices for atom interferometry [17], or even to implement elementary forms of quantum error correction [10]. For an experimentalist, the challenge thus presents itself of trying first to validate theoretical models of conditional quantum dynamics, and then to demonstrate some of these intriguing proposed applications.

As of the summer of 1995, when I first started to become aware of the fascinating physics involved in the continuous observation of quantum dynamics, it was clear that most of the basic technical tools required for such research had yet to be developed. So the essential motivation for most the work presented in this thesis has been to create both the experimental and theoretical foundations that I think will be needed, working within the concrete physical paradigm of cavity quantum electrodynamics (cavity QED). On the experimental side, this has primarily involved the construction of the first laboratory apparatus for achieving strong coupling between laser-cooled atoms and quantized electromagnetic fields in an optical cavity. Beyond initial demonstrations of this new physical system, I have worked substantially to eliminate sources of excess technical noise that compromised the fidelity of measurements made with our first-generation apparatus. My theoretical research has largely focused on generating new ideas about what should be done in experiments on con-

ditional quantum dynamics, and hence on sharpening up the goals for what really needs to be done in the lab.

Before presenting a more explicit resume of the results contained in this thesis, let me turn in the next few pages to a brief overview of the substantive (as opposed to ideological) focus of my recent work, cavity QED with laser-cooled atoms.

1.2 Methods: cavity QED with cold atoms

The basic idea in modern cavity QED is to investigate the dramatic alterations of atomic radiative processes within a high-finesse electromagnetic resonator, and to utilize them for fundamental research in fields such as quantum measurement and quantum information. The particular laboratory system of interest for this thesis will be a single ground-state (as opposed to Rydberg) ^{133}Cs atom coupled to one eigenmode of an optical Fabry-Perot resonator [18, 19, 20, 21].

The dynamics of such an atom-cavity system may be parametrized by a set of three fundamental rates: the single-photon Rabi frequency $2g_0$, the atomic dipole decay rate $\gamma_\perp$, and the cavity field decay rate κ. We require $g_0 \gg (\gamma_\perp, \kappa)$ in order for the coherent, quantum-mechanical interaction between atom and cavity to dominate the dissipative processes of spontaneous emission and cavity decay. If this all-important condition—known as *strong coupling*—can be achieved, we may conduct experiments in a regime where the coupled atom-cavity system behaves essentially as a *molecule* [22, 8]. It develops new eigenstates, which are entangled states of the atom and cavity, and it is the corresponding excitation spectrum that must be considered when computing the optical response of the atom-cavity system to a driving laser.

Experimentally, $\gamma_\perp$ is fixed once we have chosen an atom to work with, so in order to achieve strong coupling we need to find some way of making g_0 as large as possible while keeping κ relatively small. Formally, the single-photon Rabi frequency is given by $2g_0 = 2d \cdot E_1/\hbar$, where d is the atomic dipole-transition matrix element and E_1 is the intracavity electric field per photon. We can infer the general determinants of

the latter quantity by considering the equation,

$$\hbar\omega \propto \int dV \left(|E_1|^2 + |B_1|^2\right). \tag{1.7}$$

Here the integral is over the volume of space in which the photon is confined, $\hbar\omega$ is the photon energy, and the integrand is just the Maxwell energy density. Clearly when the volume of integration gets smaller and smaller, the electric and magnetic fields per photon must increase in order for the equation still to hold. Hence we meet the first condition of strong coupling by making an optical resonator with very small mode volume. A Fabry-Perot cavity consists of two spherical mirrors, so we minimize the mode volume by choosing short radii of curvature for the mirrors and by making the overall cavity length (distance between the two mirrors) as short as we can.

Of course, as we make the cavity shorter we increase the frequency with which a circulating photon has to bounce off of each mirror. In an empty cavity losses only occur during bounces, so increasing g_0 in a Fabry-Perot cavity must come at the price of also increasing κ. We can still maintain an acceptably large ratio g_0/κ, however, by using mirrors of very high reflectivity. By using mirrors with an overall transmission and loss of $\sim 1.5 \times 10^{-5}$, for example, we have $g_0/\kappa \simeq 3.4$ and $g_0/\kappa \simeq 4$ in a cavity of length ~ 107 μm that was used in the experiment to be described in this thesis. Other efforts in our group have realized ratios as high as $g_0/\kappa \simeq 18$ [23] and $g_0/\gamma_\perp \simeq 46$ [45]. For a more quantitative discussion of the art of Fabry-Perot construction, see Quentin's thesis [18].

The phenomenology of strong coupling is perhaps best characterized by a complementary pair of dimensionless quantities, the critical photon number m_0 and the critical atom number N_0 [22]:

$$m_0 \equiv \frac{4\gamma_\perp^2}{3g_0^2}, \qquad N_0 \equiv \frac{2\kappa\gamma_\perp}{g_0^2}. \tag{1.8}$$

In the strong coupling regime, both of these quantities are much less than one. Having a critical photon number below unity means that one can investigate nonlinear optics

in the quantized domain, with an average photon number less than one per mode [24]. Such capabilities formed the basis of our group's recent work on demonstrating conditional phase shifts for quantum logic [25].

My experimental thesis research, on the other hand, has uniquely been enabled by having a critical atom number $N_0 \ll 1$. If we think about the optical cavity as an input-output device (in the sense that you can shine a laser in one end and see what comes out the back), the critical atom is a measure of how many atoms we have to put inside the cavity before its optical response is drastically changed. One might ask what the use is of having $N_0 \ll 1$, as atoms certainly come only in integral units. Roughly speaking, the smaller N_0 is, the greater the sensitivity and bandwidth we have for monitoring changes in the *effective* number of intracavity atoms. The actual atom-cavity coupling strength $g(\vec{\mathbf{r}})$ should clearly vary with the atom's position inside the cavity, and for Cesium it also happens to vary fairly strongly with the atomic internal (Zeeman) state.

The enhanced coupling between atoms and photons inside a high-finesse optical cavity thus provides a novel basis for optical measurements that continuously monitor atomic degrees of freedom. Experimentally, the principle challenge to realizing such measurements is the need to arrange a sufficient separation of timescales between the achievable measurement bandwidths and the evolution of monitored atomic variables. This is not an altogether trivial constraint, and in fact no cavity QED experiment prior to our Optics Letter of 1996 [26] had managed to observe single-atom effects on cavity transmission without integrating over the transits of many successive atoms through a cavity. In that publication, we demonstrated that the use of simple *laser cooling* techniques allowed one to achieve the required separation of timescales, and that we did indeed thus gain the ability to detect the passage of single Cesium atoms through an optical cavity *in real time*, and with high signal-to-noise ratio.

To perform single atom detection, we monitored the transmission of a resonant cavity probe laser after dropping atoms from a magnetooptic trap (MOT) [27] formed over the gap between the cavity mirrors (see Figure 1.1). Before the arrival of the freely-falling atoms, the cavity transmits a constant mean power P_0 which is deter-

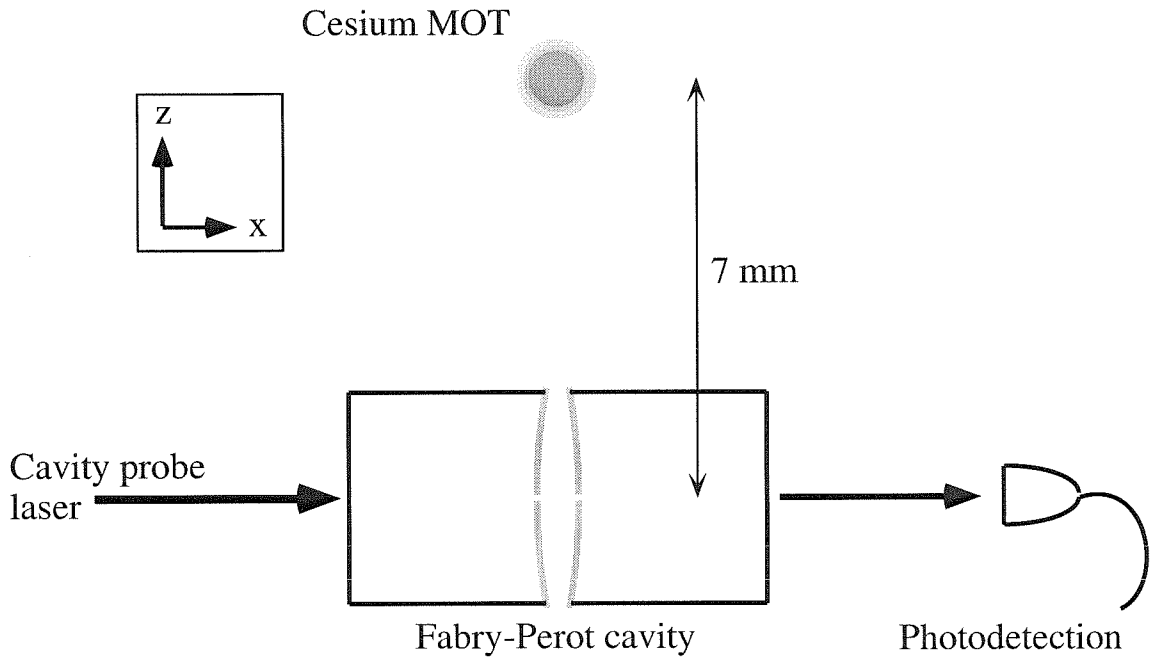

Figure 1.1: Schematic diagram of the basic experiment.

mined with a balanced heterodyne detector. After a mean time delay corresponding to that required for atoms to fall from the MOT to the cavity axis, the transmitted power exhibits distinct dips associated with intracavity absorption by single atoms falling through the resonator mode volume.

In Figure 1.2 we display several examples of the time-varying heterodyne signal recorded after each dropping of the MOT. Each of these is representative of the traces that can be seen in real time, on an analog oscilloscope, while the experiment is running. There is clearly a large dispersion in the magnitude of the atom-transit signals, just as one would expect since some atoms should fall through the wings of the cavity mode volume while others fall right through the center. Note the logarithmic vertical scale on these plots—the directly recorded data has units of dBm (photocurrent power in the resolution bandwidth) because the heterodyne signal was processed and detected with an HP Spectrum Analyzer (see below). Thus, some of the atom-transit signals exhibit very high contrast ($\sim$ 100:1), but one should bear in mind that the trace "width" associated with shot noise and technical noise is not constant along the vertical scale of the graph.

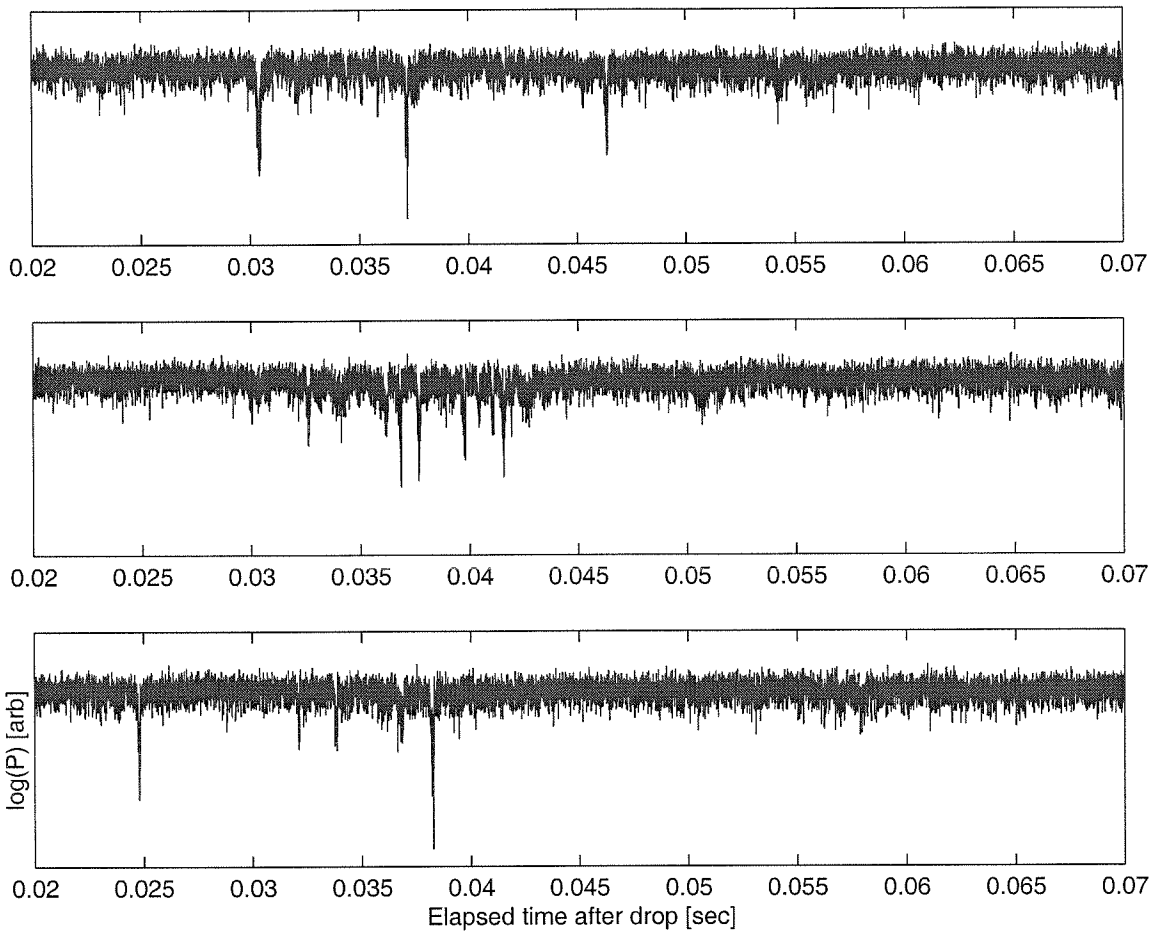

Figure 1.2: Time-varying cavity transmission after dropping a cesium MOT.

Figure 1.3 shows a histogram of atom arrival times for a typical run of 450 trap-drop cycles. In this run we found an average 2.5 events/drop, although the frequency of events can be made as high as $\sim$ 50 prominent events/drop by optimizing the size of the MOT. The mean arrival time in the histogram is just what it should be for atoms falling 7 mm under the influence of gravity (this distance was verified independently), and the variance ($\pm$7 ms) of the arrival-time distribution is as expected for an atomic sample cooled to the Doppler limit. We also show a similar histogram, generated by the same automatic event-finding routine from a data set taken with the MOT beams blocked. The marked difference between the two histograms demonstrates that we have very high signal-to-noise for identifying atomic transit signals in the heterodyne photocurrent.

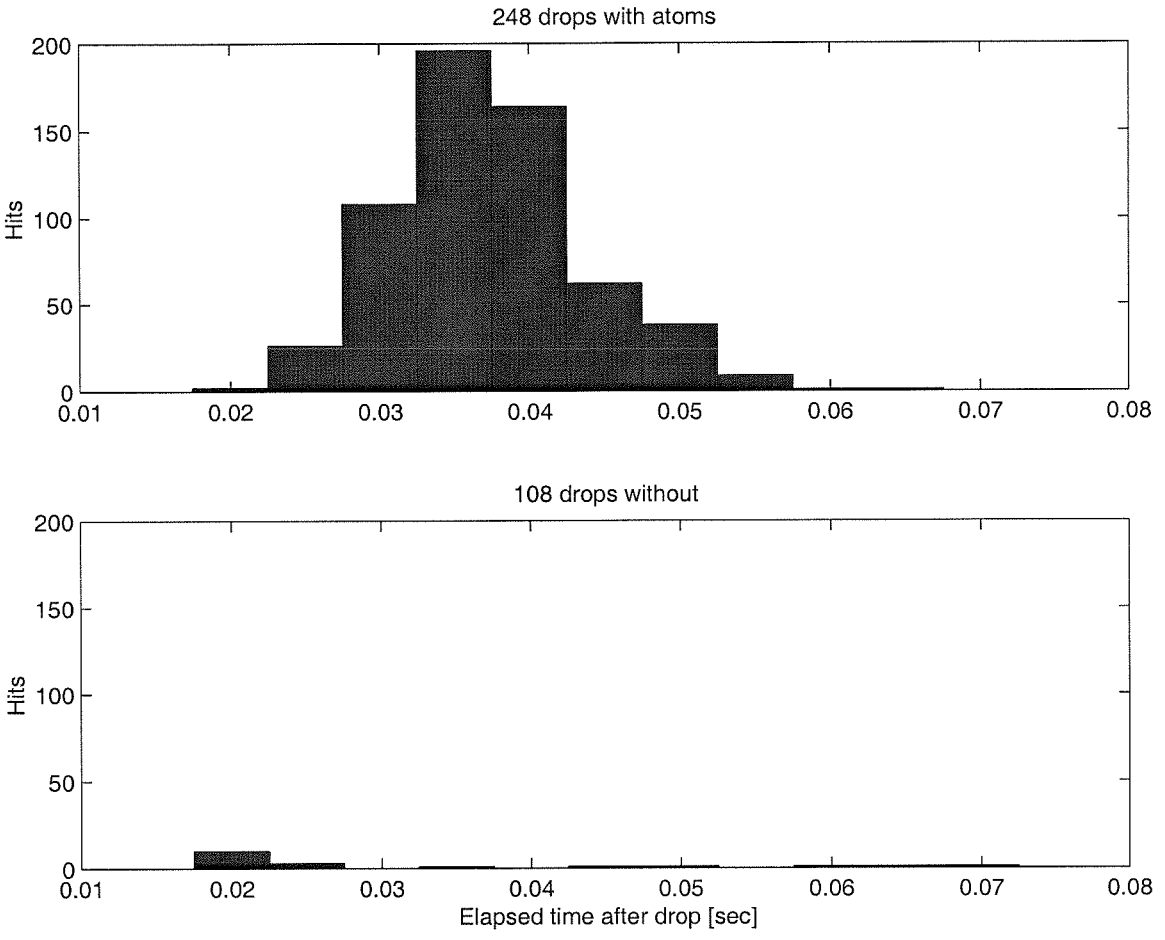

Figure 1.3: Histograms of "atom" arrival times (see text).

What really got us excited about the 1996 data, however, was what we saw when we took individual single-atom transit signals and viewed them on an expanded time scale. Figure 1.4 shows six notable examples that we put into the Optics Letter. I had initially expected that we would see a smooth, gaussian envelope for the individual atom transits, but instead we saw quite a few examples with lots of internal structure. After thinking about it a bit, it seemed possible that we might be seeing some signature of atomic motion relative to the *standing-wave* structure of the cavity eigenmode. Since the two mirrors of the Fabry-Perot cavity resonator present reflecting boundary conditions for the intracavity field, the distribution of electric field between them is a standing wave and correspondingly has nodes and anti-nodes at a spacing of one-fourth the optical wavelength. If the atom happens to drift onto

a node while it falls through the cavity, it temporarily decouples from the field. So at those moments in time, we would expect the cavity transmission to come all the way back up to the empty-cavity level. From the geometry of the cavity and MOT, one can estimate that the average velocity of an atom along the cavity standing wave should be $\sim$ 1 cm/s as it enters the cavity mode volume. Again from geometry, we inferred that it should therefore coast over five or six periods of the standing wave within its vertical transit time. That seemed consistent with the frequency of the oscillations seen within individual transit signals, although it was also clear that we were bandwidth-limited. The overall width of each transit signal is approximately 250 μs, and the timescale for the major features was around $20 - 30$ μs. Our digitization rate back in those days was only 500 kHz.

If it were really true that we could resolve atomic motion relative to the intracavity standing wave, over timescales on the order of 100 μs, we thought that we should really be at the Standard Quantum Limit for monitoring the position of a free mass [28, 5]. Taking the standard expression,

$$\Delta x_{SQL} = \sqrt{\frac{\hbar\tau}{2m}} \approx 150\text{nm}, \tag{1.9}$$

where $\tau = 100$ μs is the measurement "interval" and m is the particle (here atomic) mass. The distance associated with oscillations in the signal due to motion over the standing wave, for comparison, would be $\lambda/4 \simeq 213$ nm. One should be careful to note, however, that the above expression is really derived for a scenario with two *discrete* measurements of the atom's position, separated in time by an interval τ. Our measurements, on the other hand, seemed to be nearly continuous in nature. So we weren't exactly sure what to say about the potential significance of our measurements, but it certainly seemed imperative to get back in the lab and clarify what we were seeing. Dimensionally, Δx_{SQL} is the only fundamental length scale you can come up with for quantum measurements of position, so we had faith that something interesting must happen if you could get down to that level.

In terms of the experiment, it was immediately clear that two things had to hap-

pen. First, we had to obtain much higher analogue bandwidth and digital sampling rate to gain some reserve between the standing-wave oscillations and the Nyquist frequency. Second, we needed to find a better way to stabilize the cavity length, because our initial scheme required us to dither the cavity length at a frequency $\sim 80 - 100$ kHz in order to obtain a lock-in error signal. It was not inconceivable that this dither was responsible for at least some of the structure we saw in our early data. Soon after we set about these two tasks, however, it also became clear that had a more fundamental problem in the form of *heating* of the atomic motion. Because we drove the cavity with a probe beam at the atomic resonance frequency, any given atom would absorb and emit roughly 2000 photons during the time it took to fall through the cavity. Hence, even if it entered the top of the mode volume with a velocity ~ 1 cm/s along the standing wave, it could quickly be kicked up to 1 m/s or greater. The real severity of this problem was pointed out to us by Andrew Doherty, a theory student of Dan Walls at the University of Auckland, who visited our group for several months in order to initiate collaborative work that continues to this day. When the atomic velocities along the standing wave exceeded even 10 cm/s, the resulting oscillations in the heterodyne signal moved above the passband of the anti-aliasing filters we had to apply to the photocurrent [29].

Prior to our experimental results in 1996, Quadt, Collett, and Walls [30] had published a numerical simulation in which they showed that one should be able to track atomic motion in a cavity by detuning the probe beam far from atomic resonance, and by measuring the *phase* of the light transmitted through the cavity. Encouraged by this result, we decided that we would just have to learn to make high-bandwidth, shot-noise-limited optical phase measurements and repeat the falling atoms experiment with a far-detuned probe beam. This was to be done in conjunction with some general upgrade of the apparatus in terms of our signal bandwidth and the cavity locking scheme. Little did we know what we were getting ourselves into.

In addition to these experimental travails, there was the unresolved interpretational matter of what exactly the quantity Δx_{SQL} was supposed to signify in the context of "continuous" position measurements. Some theoretical work had been

done for the special case of *narrowband* AC position measurements on a damped harmonic oscillator well above its natural frequency, but we were really aiming at broadband measurements with sensitivity over nearly two decades of frequency (from the inverse transit time $\sim$ 4 kHz to our eventual 6 dB bandwidth of 300 kHz). So it seemed that we needed to spend some time thinking about quantum measurement theory as it applied to the experiment, in addition to working on technical upgrades in the lab.

That was the general situation two years ago, which brings us basically up to the period that I have documented in this thesis.

1.3 Results: continuous observation of single-atom dynamics

As of the writing of this thesis, I am fairly well satisfied that the outstanding issues raised by our initial data of 1996 have now been resolved. Equation (1.9) is *not* the right way to think about quantum limits to broadband position measurement, and I believe that the real story is as described in the first section of Chapter 4. My essential conclusion from theoretical work is that one really needs to approach the experiment with a sincere ambition to do *precision* measurement, in the sense that nothing much of real scientific interest can be said about the data unless we manage to suppress all sources of excess technical noise and reach the fundamental quantum limits of measurement sensitivity.

In our actual experiment we do now reach the fundamental sensitivity set by heterodyne shot noise, although a further factor of $\sqrt{3}$ could in principle be gained by improving the photodetection efficiency. Chapter 2 provides a detailed technical evaluation of the substantial improvements we have made on the experimental apparatus, and discusses the new data that we have obtained for monitoring atom transits with a detuned probe beam. I also discuss my final views on resolving atomic motion relative to the intracavity standing wave, which are essentially that we now have the sensitiv-

ity and bandwidth to do so but suffer from a sort of "confusion limit" in which the signals we need to see are buried in background contributions from transverse atomic motion and/or optical pumping among atomic Zeeman states. Chapter 3 fills in some technical details that were omitted from Chapter 2, and also discusses some of the things that we tried to do early on and the reasons why we had to abandon them.

Getting back to the subject discussed in the first section of this chapter, I now consider the real significance of Standard Quantum Limits in broadband measurement scenarios to be the fact that they set a concrete benchmark that must be reached *before* one can really talk about doing experiments on conditional quantum dynamics. It seems that we have very nearly reached such a benchmark in our most recent experimental work, so in section 2 of Chapter 4 and in Appendix B I have included some thoughts on the kinds of things we should be setting our sights on for the next round of experimentation.

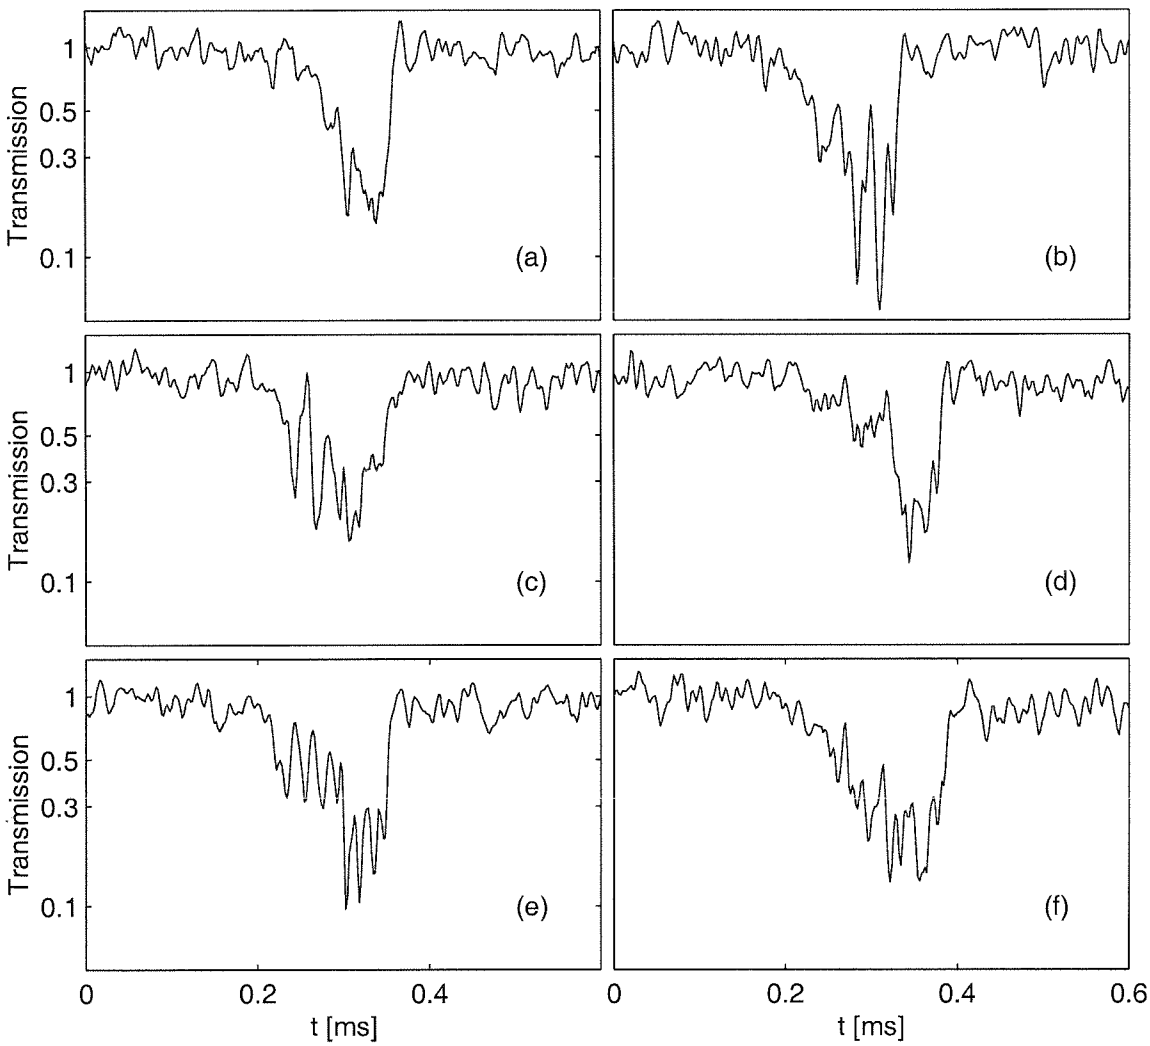

Figure 1.4: Six examples of transit signals associated with the passage of individual atoms through the optical cavity. The data are normalized and taken with the same parameters as Figure 1.2 except that P_0 corresponds to 3 intracavity photons.

Chapter 2 Full observation of single-atom dynamics in cavity QED

Optical cavity quantum electrodynamics (QED) in the strong coupling regime [22] provides a truly unique experimental paradigm for real-time observation of quantum dynamical processes at the *single-atom* level. Spectacular advances have certainly been made in the preparation and tomography of quantum states of motion for a single trapped ion [33, 34], but all such experiments have involved the gradual accumulation of ensemble-averaged data over many successive realizations of the process being studied. Recent studies of single-molecule dynamics have likewise demonstrated the "immediate" detection of photochemical [31] or conformational [32] events, but I would stress that such experiments presently lack the potential that cavity QED provides for observing quantum processes on a timescale that makes coherent control/intervention a tangible possibility.

Real-time observation of quantum dynamics has recently been achieved in *many*-atom systems, in experiments on the vibrational excitations of a trapped Bose-Einstein condensate [35] and on the decay of coherent oscillations of an ensemble of atoms in an optical lattice [36, 37]. In contrast to programs like these, for which the scientific emphasis lies mainly on noninvasive observation of a system's intrinsic dynamical processes, experiments in single-atom cavity QED hold great potential for enabling the development and validation of a quantitative understanding of exactly how measurement backaction *alters* the dynamical behavior of a continuously-observed open quantum system [38, 4]. A theoretical basis for this endeavor is presently maturing in the context of quantum trajectory formalisms [8, 7, 39, 13], but some significant technical challenges remain to be solved before definitive experiments can be performed in the lab. In the past two years we have managed to make substantial progress on this front, so I hope that it will not be too much longer before

we can begin to perform quantitative tests of measurement-based stochastic Master equations, and perhaps even to implement some recently-proposed "applications" for the continuous observation of dissipative quantum dynamics [40, 10, 41, 42, 14, 17].

In this chapter I focus on a detailed description of recent experiments in which we have used broadband heterodyne spectroscopy to record the complete time-evolution of interaction energy between one atom and a high-finesse optical cavity, during individual scattering events of ~ 250 μs duration. The instantaneous value of this energy depends on both the atomic position and (Zeeman) internal state, with variations in time caused predominantly by atomic motion through the spatial structure of the cavity eigenmode. Our measurements have been conducted in a regime of strong but dispersive atom-cavity coupling, ensuring that dynamical variations of the interaction energy remain within the shot-noise-limited detection bandwidth. With characteristic atom-cavity interaction energies $E_{\rm int}/\hbar \sim 10$ MHz, we achieve measurement sensitivities $\simeq 3$ kHz/$\sqrt{\rm Hz}$ over a bandwidth that covers the dominant rates of variation in $E_{\rm int}$.

Unlike typical pump-probe measurements of scattering dynamics in real (*e.g.*, diatomic) molecular systems [43, 44], our experiments on the Jaynes-Cummings molecule yield a continuous time-domain record of the atom-cavity coupling during each individual scattering event. The data clearly illustrate variations caused by atomic motion through the spatial structure of the cavity eigenmode and/or optical pumping among the atomic internal (Zeeman) states. In certain parameter regimes of the detuning and cavity driving strength, distinctive indications of the quantum-mechanical nature of the atom-cavity coupling can be seen in the photocurrent recorded from just a *single* atomic transit.

Because of the standing-wave spatial structure of the cavity eigenmode and the corresponding rapid variation of the atom-cavity coupling strength over sub-wavelength distances, our data should display a *sensitivity* of 10^{-10} m/$\sqrt{\rm Hz}$ to atomic displacements along the cavity axis. Unfortunately, we cannot realize this figure as a *precision* for monitoring the atomic position, as we do not presently have any means of separating signal variations due to motion through the standing wave from "background"

contributions due to transverse motion or optical pumping. In the final section of this chapter, I shall discuss the motivation for further work to disambiguate the nature of rapid variations in the observed signals.

2.1 Basic theoretical description

In simple terms, our experimental procedure is to drop a cloud of cold Cesium atoms from a magnetooptic trap (MOT) into a high-finesse optical cavity, while continuously monitoring the cavity's complex susceptibility with a weak probe laser [26, 45]. By limiting the number of atoms in the initial cloud, we can easily reach an operating regime in which atoms transit the cavity only one at a time. Using broadband heterodyne detection and a high-speed digitizer, we continuously record both the amplitude and phase of the transmitted probe beam during ~ 50 ms time windows. Each window typically contains from zero to five atom transit signals.

The elementary theoretical description of such a measurement employs steady-state solutions of the nonselective Master Equation for a *stationary* two-level atom coupled to a single electromagnetic mode via the Jaynes-Cummings interaction Hamiltonian. In this treatment, optical pumping among Zeeman states and atomic center-of-mass motion can only be included at the level of a time-dependent atom-cavity coupling strength. Although this type of approach cannot make predictions about dynamical variations in the coupling strength, it provides at least a "first-order" basis for interpreting some time-independent features of our data.

If we let ρ denote the density operator for the joint state of the atom and cavity, the nonselective Master equation (in the electric dipole and rotating-wave approximations) reads:

$$\begin{aligned}
\dot{\rho} &= \frac{-i}{\hbar}\left[\hat{H}_0, \rho\right] + \gamma_\perp\left(2\hat{\sigma}\rho\hat{\sigma}^\dagger - \hat{\sigma}^\dagger\hat{\sigma}\rho - \rho\hat{\sigma}^\dagger\hat{\sigma}\right) \\
&\quad + \left(\kappa_a + \kappa_b + \kappa_c\right)\left(2\hat{a}\rho\hat{a}^\dagger - \hat{a}^\dagger\hat{a}\rho - \hat{a}^\dagger\hat{a}\rho\right), \qquad (2.1) \\
\hat{H}_0 &= \hbar\Theta\hat{a}^\dagger\hat{a} + \hbar\Delta\hat{\sigma}^\dagger\hat{\sigma} + \sqrt{2\kappa_a}\mathcal{E}\left(\hat{a} + \hat{a}^\dagger\right) + H_{\rm int}, \qquad (2.2) \\
H_{\rm int} &= \hbar g_0 e^{-\left(y^2+z^2\right)/w^2}\cos\left(k_L x\right)\left[\hat{a}\hat{\sigma}^\dagger + \hat{a}^\dagger\hat{\sigma}\right]. \qquad (2.3)
\end{aligned}$$

Here $\gamma_\perp$ is the atomic dipole decay rate, κ_a is the cavity field decay rate through the input mirror, κ_b is the cavity field decay rate through the output mirror, κ_c is the cavity field decay rate due to intracavity scattering/absorption losses, $\Delta \equiv \nu_a - \nu_p$ is the atom-probe detuning, $\Theta \equiv \nu_c - \nu_p$ is the cavity-probe detuning, and the coupling strength g_0 is equal to half the maximum single-photon Rabi frequency. We treat the atomic center-of-mass coordinates x, y, z as c-number parameters, with the x-axis coinciding with the cavity axis and z parallel to gravity. We have written the Master Equation in a frame rotating at the drive frequency, so $\mathcal{E}$ is a constant term proportional to the complex amplitude of the driving field.

To find the steady-state density operator as a function of driving strength and various detunings, we simply set $\dot{\rho}_{ss} = 0$ and solve for ρ_{ss} using linear algebra. The expected amplitude and phase of the heterodyne photocurrent may then be computed as [46, 7]

$$\langle i_{het}(t) \rangle = \eta f_L^{1/2} \sqrt{2\kappa_b} \exp\left[i\left(\Omega_L t + \phi_L\right)\right] \mathrm{Tr}\left[\rho_{ss}\hat{a}\right], \tag{2.4}$$

where η represents the overall photodetection efficiency (including propagation losses between the cavity and photodetectors, heterodyne efficiency, and detector quantum efficiency), f_L and ϕ_L represent the photon flux and phase of the (optical) local oscillator, and Ω_L is the frequency of the optical local oscillator relative to the rotating frame (cavity driving field). In the experiment, we mix $i_{het}(t)$ with an rf local oscillator at the frequency Ω_L (which ranges between 40 and 190 MHz) and separately record the in-phase and quadrature components of the slowly-varying envelope (with an analogue bandwidth of 300 kHz).

This system can also be treated semiclassically, using the optical bistability state equation (OBSE) [47]. The OBSE is traditionally written in terms of the scaled field variables x and y (not to be confused with the atomic coordinates), with the correspondence

$$x \equiv \frac{\langle \hat{a} \rangle}{\sqrt{m_0}}, \quad y \equiv \frac{\sqrt{2\kappa_a}\mathcal{E}}{(\kappa_a + \kappa_b)\sqrt{m_0}}, \tag{2.5}$$

where the saturation photon number m_0 is given (for a gaussian, standing-wave cavity mode) by $\frac{4\gamma_\perp^2}{3g_0^2}$. For a given driving strength $\mathcal{E}$, the expected intracavity field amplitude

can be found by inverting the equation

$$y = x\left(1 + \frac{2C}{1+\delta^2+x^2} + i\frac{\phi - 2C\delta}{1+\delta^2+x^2}\right), \tag{2.6}$$

where we again work in a rotating frame at the drive frequency, $\delta \equiv (\omega_a - \omega_p)/\gamma_\perp$ represents the scaled atom-probe detuning, $\phi \equiv (\omega_c - \omega_p)/(\kappa_a + \kappa_b + \kappa_c)$ represents the scaled cavity-probe detuning, and the "cooperativity" C is defined by

$$C \equiv \frac{g^2}{2(\kappa_a + \kappa_b + \kappa_c)\gamma_\perp}. \tag{2.7}$$

Note that the dependence of g here on atomic position and internal state is implicit. A semiclassical prediction for the heterodyne photocurrent is obtained by substituting $x\sqrt{m_0}$ for $\mathrm{Tr}\left[\rho_{ss}\hat{a}\right]$ in equation (2.4).

Although I won't actually make use of the equation below, I would like to note that it should be possible to write a comprehensive quantum theory for the measurements we perform. The selective evolution of the atom-cavity state should follow a Stochastic Master Equation (SME) for continuous heterodyne measurement of the cavity leakage field, as described in [7, 39]. To account for departures from the ideal scenario of a two-level atom and single field mode, we should expand the Hamiltonian to account for the full manifold of Zeeman states in the Cesium D2 ($F = 4 \rightarrow F' = 5$) system, quantize the atomic center-of-mass degrees of freedom, and deal with non-degeneracy of polarization eigenmodes in the cavity. In its un-normalized form, I believe that such an SME should read

$$\begin{aligned}
\frac{d}{dt}\tilde{\rho}_c &= \frac{-i}{\hbar}\left[\hat{H}_0, \tilde{\rho}_c\right] - \gamma_\perp\left(\hat{\Sigma}^\dagger\hat{\Sigma}_{tot}\rho_c - \rho_c\hat{\Sigma}^\dagger_{tot}\hat{\Sigma}_{tot}\right) \\
&+ 2\gamma_\perp\sum_{\varepsilon=0,\pm1}\int d\mathbf{u}\,\mathcal{N}_\varepsilon(\mathbf{u})e^{-ik_L\mathbf{u}\cdot\hat{\mathbf{r}}}\hat{\Sigma}_\varepsilon\tilde{\rho}_c\hat{\Sigma}^\dagger_\varepsilon e^{ik_L\mathbf{u}\cdot\hat{\mathbf{r}}} \\
&+ \kappa_2\left(2\hat{a}_2\tilde{\rho}_c\hat{a}_2^\dagger - \hat{a}_2^\dagger\hat{a}_2\tilde{\rho}_c - \tilde{\rho}_c\hat{a}_2^\dagger\hat{a}_2\right) \\
&+ \kappa_1\left(2\langle\hat{a}_1\rangle_c\tilde{\rho}_c\hat{a}_1^\dagger + 2\langle\hat{a}_1^\dagger\rangle_c\hat{a}_1\tilde{\rho}_c - \hat{a}_1^\dagger\hat{a}_1\tilde{\rho}_c - \tilde{\rho}_c\hat{a}_1^\dagger\hat{a}_1\right) \\
&+ \sqrt{2\kappa_1}\left(\tilde{\rho}_c\hat{a}_1^\dagger\xi^* - \hat{a}_1\tilde{\rho}_c\xi\right) + 2\kappa_1\hat{a}_1\tilde{\rho}_c\hat{a}_1^\dagger|\xi|^2,
\end{aligned} \tag{2.8}$$

$$\begin{aligned}
\hat{H}_0 &= \frac{\hat{p}^2}{2m} + \hbar\Theta\hat{\Pi}_e + \mathbf{B}\cdot\left(\mu_g\hat{\mathbf{J}}_\mathbf{g} + \mu_e\hat{\mathbf{J}}_\mathbf{e}\right) \\
&+ \hbar\Delta_1\hat{a}_1^\dagger\hat{a}_1 + \mathcal{E}_1\left(\hat{a}_1 + \hat{a}_1^\dagger\right) \\
&+ \hbar\Delta_2\hat{a}_2^\dagger\hat{a}_2 + \mathcal{E}_2\left(\hat{a}_2 + \hat{a}_2^\dagger\right) + H_{\text{int}}, \qquad (2.9) \\
\hat{H}_{\text{int}} &= \hbar g_0 e^{-\left(\hat{y}^2+\hat{z}^2\right)/w^2}\cos\left(k_L\hat{x}\right)\left[\hat{a}_-\hat{\Sigma}_{-1}^\dagger + \hat{a}_-^\dagger\hat{\Sigma}_{-1}\right. \\
&+ \left.\hat{a}_+\hat{\Sigma}_{+1}^\dagger + \hat{a}_+^\dagger\hat{\Sigma}_{+1}\right], \qquad (2.10) \\
\hat{\Sigma}_\varepsilon &\equiv \sum_{m=-4}^{4} |g,m\rangle\langle e,m+\varepsilon|\cdot\mathcal{C}\left(4,m|1,\varepsilon|5,m+\varepsilon\right), \qquad (2.11) \\
\hat{\Sigma}_{tot} &\equiv \sum_{\varepsilon=0,\pm}\hat{\Sigma}_\varepsilon. \qquad (2.12)
\end{aligned}$$

Note that I have not maintained a distinction between cavity losses associated with the input mirror, output mirror, and intracavity absorption/scattering. Here ξdt is a complex Wiener increment [7], $\mathcal{N}_\epsilon(\mathbf{u})$ denotes the angular distribution for ε-polarized spontaneous emission [48], $\hat{\Pi}_e$ is a projector into the atomic excited-state manifold, $\mu_g\hat{\mathbf{J}}_\mathbf{g}$ and $\mu_e\hat{\mathbf{J}}_\mathbf{e}$ are total magnetic-moment operators for the atomic ground and excited levels, $\mathbf{B}$ is the ambient magnetic field, and the $\mathcal{C}$ are relative weights for the various Zeeman transitions [49]. We account for cavity birefringence by defining

$$\begin{pmatrix}\hat{a}_-\\ \hat{a}_+\end{pmatrix} \equiv \begin{pmatrix}c_1^- & c_2^-\\ c_1^+ & c_2^+\end{pmatrix}\begin{pmatrix}\hat{a}_1\\ \hat{a}_2\end{pmatrix}, \qquad (2.13)$$

where the matrix of coefficients is the transpose of that which defines the transformation from the circular-polarization unit vectors (in the y,z-plane) to the cavity eigenpolarizations. The driving fields $\mathcal{E}_{1,2}$ are similarly defined. One further refinement should be made to account for imperfect photodetection efficiency, but I refer the reader to reference [7] for details.

2.2 Experimental apparatus and procedures

Figure 2.1 provides a general overview of the apparatus, indicating the schematic arrangement of various components to be described below. The diode laser setup for

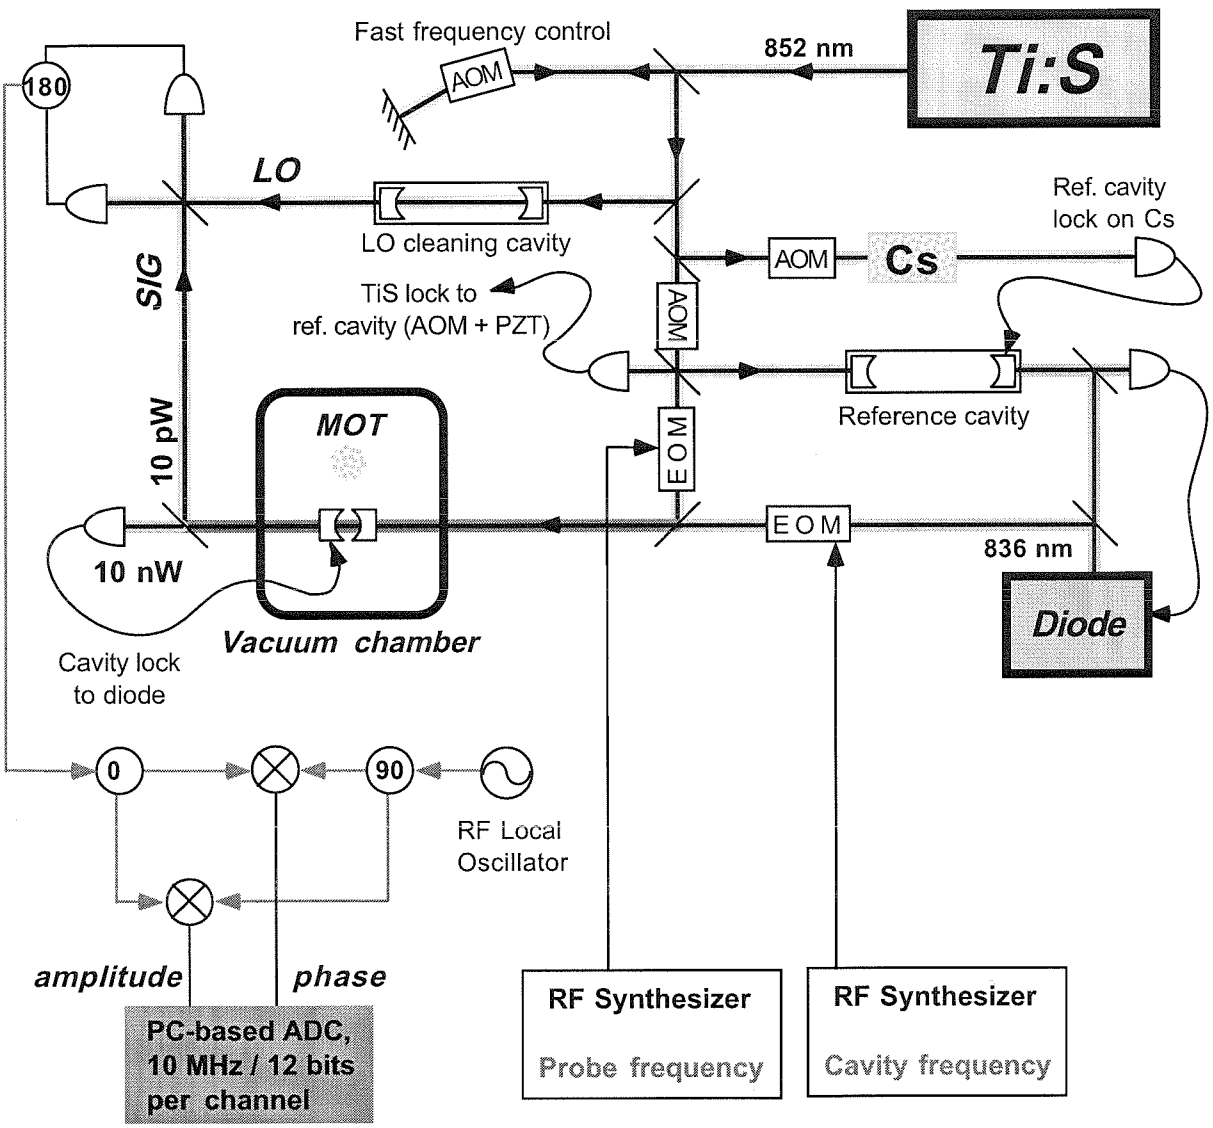

Figure 2.1: Schematic overview of the apparatus.

forming the Cs MOT is not shown.

2.2.1 High-finesse microcavity

We use a Fabry-Perot high-finesse microcavity ("physics cavity") consisting of two spherical mirrors with 1 m radius of curvature [50]. The cavity was constructed with a mean length $l \simeq 107.5$ μm, which we inferred from the cavity's measured free spectral range of 1.395×10^{12} Hz. The measured l and specified radii of curvature geometrically determine the cavity's electromagnetic mode volume for TEM_{00} modes near 852 nm [51, 18]. Together with the dipole decay rate $\gamma_\perp/2\pi \simeq 2.6$ MHz for the Cs $6P_{3/2}$ level [52], this determines our optimal coupling constant $g_0/2\pi$ to be $\simeq 11$ MHz for $\sigma_\pm$ transistions (specifically the $6S_{1/2}\,(F = 4, m_F = \pm 4) \rightarrow 6P_{3/2}\,(F = 5, m_F = \pm 5)$) and $\simeq 6$ MHz for π transitions ($6S_{1/2}\,(F = 4, m_F = 0) \rightarrow 6P_{3/2}\,(F = 5, m_F = 0)$) within the D2 Zeeman manifold [22].

In order to maintain a reasonable opening between the edges of the mirror substrates with this small l, we asked the mirror manufacturer to reduce the substrate diameters from the standard value of 7.75 mm down to 3 mm. This machining was done *after* the substrates had been superpolished and coated, but this process did not seem significantly to degrade the mirror quality. The nominal combined transmission and loss per mirror, before machining, was 1.5×10^{-5}. Direct measurements of the cavity finesse yield $\mathcal{F} \simeq 217,000$ at an optical wavelength of 852.36 nm, consistent with a combined mirror transmission and loss $T^* \simeq 1.45 \times 10^{-5}$. This value of T^* is inferred from the measured value of l and the measured cavity HWHM $\kappa/2\pi \simeq 3.21$ MHz.

The cavity used throughout the work described in this paper suffers from a rather pronounced birefringence, which for TEM_{00} modes near 852 nm induced a splitting of 8 ± 2 MHz between linearly-polarized eigenmodes. It is not entirely clear whether this birefrigence is a result of the substrate machining for diameter-reduction, a property of the coatings, or something associated with the mirror-mounting procedure used for this particular cavity. We note that recent efforts by other members of our group

[45] have produced a cavity of approximately the same finesse with greatly-reduced birefringence, with mirrors from a different coating run and with great care taken to minimize cavity misalignments and stress on the mirror substrates.

Our two mirror substrates are mounted in vee-grooves atop independent aluminum blocks, with a piezoelectric actuator between the blocks for active servo-control of the mirror separation (cavity length). The cavity mount sits on a stack of alternating OFHC copper blocks and viton o-rings for passive vibration isolation, all within an ion-pumped vacuum chamber whose background pressure was typically $\sim 10^{-8}$ torr (inferred from the ion pump current). In order to bring the MOT as close as possible to the central axis of the physics cavity, we had to use a rather open (and non-magnetic) mount design, leading to some compromises in the way of mechanical stability. Sitting on the vibration-isolation stack and under vacuum, we found that the native noise spectrum of the cavity length extended out to about 4 kHz, with one prominent resonance at 50 Hz (which we attribute to a transmission resonance of the isolation stack). Due to a set of PZT-actuator resonances above 10 kHz, we have ultimately been limited to a unity-gain bandwidth $\simeq 1$ kHz for the cavity stabilization servo.

2.2.2 Laser and cavity locking schemes

Although the principle aim for this experiment is to stabilize the cavity length at some precise offset $(\Delta - \Theta)/2\pi \sim 0 - 100$ MHz from the Cs D2 resonance at 852.359 nm, the strong atom-cavity coupling places severe restrictions on the optical power that can be used for the purpose of generating an error signal. On resonance, the saturation intracavity photon number for our cavity is $m_0 \equiv 3\gamma_\perp^2/4g_0^2 \simeq 0.1$, which sets a fiducial cavity throughput of $4\pi\kappa_b m_0 \sim 1$ pW. With such low optical power it would be extremely difficult to obtain a high-quality error signal for locking the physics cavity. Other experiments performed in our group have circumvented this problem by using a chopped locking scheme, in which a strong "lock beam" alternates with a weak "probe beam" at 50% duty cycle and $\sim 1 - 4$ kHz frequency [24, 26, 45]. However, such a strategy inherently limits the servo unity-gain bandwidths to ~ 100

Hz at best, and would not be appropriate for future experiments with atoms trapped inside the cavity for long periods of time.

In this work we have developed an alternative locking scheme for the physics cavity, which employs an auxilliary diode laser at 836 nm to monitor the cavity length on a different longitudinal mode than that which couples strongly to intracavity Cs atoms. We use a commercial grating-stabilized diode produced by New Focus (Santa Clara, California). With a 16 nm detuning, we can send $\sim$ 10 nW through the cavity and incur an AC Stark shift of only $\sim$ 15 kHz. Using an EG&G avalanche-photodiode/transimpedance amplifier module (model C30998) for AC detection of transmitted 836 nm light, we obtain an FM error signal (modulation frequency 3.8 MHz) with signal-to-noise ratio (SNR) $\sim$ 20 at 30 kHz bandwidth.

The 836 nm diode laser is stabilized to an auxilliary "transfer cavity," which consists of a pair of 25 cm radius-of-curvature mirrors at $\simeq$ 16 cm separation. One of the mirrors is mounted on a piezoelectric actuator to allow cancellation of DC drift and low-frequency noise. The transfer cavity has a linewidth $\simeq$ 100 kHz at both 836 nm and 852 nm, with an overall mode spacing $\sim$ 300 MHz. The transfer cavity is also used for pre-stabilization of the Ti:Sapphire laser, and some of the Ti:Sapphire light is used in a Cs modulation-transfer spectrometer [53] to provide an absolute reference for feedback to the transfer cavity length. From run to run we used one or two acoustooptic modulators to offset the Ti:Sapphire/transfer-cavity lock point by +140, +87, or +43 MHz relative to the Cs $6S_{1/2}(F = 4) \rightarrow 6P_{3/2}(F' = 5)$ transition.

By comparison of the Ti:Sapphire and diode laser error signals in their respective locks to the transfer cavity, we infer that the relative rms jitter between them is $\leq$ 10 kHz. For both laser locks we use the Pound-Drever-Hall technique [54] of detecting an FM signal in reflection from the tranfser cavity. The stability of the transfer cavity resonances with respect to atomic Cs lines in a vapor cell was such that we did not see any relative jitter beyond the measurement noise in our modulation-transfer spectrometer (SNR $\sim$ 50 : 1 in 30 kHz bandwidth).

The Ti:Sapphire stabilization employs two feedback loops, one with $\sim$ 10 kHz bandwidth to the tweeter inside the ring laser and another with $\sim$ 100 kHz bandwidth

to a VCO-driven, double-passed acoustooptic modulator (AOM) just outside the laser cavity [55, 56]. The mean frequency of the AOM is 76 MHz, and the error signal going to the VCO has a lower cutoff of $\sim$ 10 kHz to prevent DC drifts. We note that use of the AOM is crucial for achieving high stability of the Ti:Sapphire frequency. The diode laser lock also uses two loops, with one providing integral gain via feedback to the grating PZT and another providing proportional gain out to $\sim$ 1 MHz via direct modulation of current going to the diode laser.

Having locked both the diode laser and Ti:Sapphire to modes of the transfer cavity, which itself is locked to Cs, we use a travelling-wave electrooptic modulator to generate an rf sideband of the diode laser at $\Delta f_0 \sim 200 - 500$ MHz. Either the upper or lower sideband is used to derive an FM error signal for locking the physics cavity by dithering Δf_0 at 3.8 MHz. This allows us to achieve arbitrary placement of the physics cavity mode near 852 nm, using the tunability of f_0.

2.2.3 Evaluation of the physics-cavity servo

Our basic requirement for the quality of the physics-cavity servo was that relative jitter of the cavity resonance and the probe laser frequency should not contribute a significant amount of noise in the heterodyne photocurrent. Hence the relevant comparison to make is between the noise in both quadratures of a demodulated beatnote and the photocurrent fluctuations produced by the local oscillator alone. Figure 2.2 shows three traces of quadrature amplitude versus time: the local oscillator alone (BOT), the phase quadrature of a beatnote produced by light transmitted through the cavity (MID), and the amplitude quadrature of light transmitted through the cavity (TOP). The data were filtered down to a bandwidth of approximately 200 kHz. At this bandwidth, the standard deviations of the phase and amplitude quadratures of the transmitted probe beam were respectively 1.01 and 1.39, relative to that of the local oscillator alone (over the 10 ms time-window shown in the figure).

Figure 2.3 shows the power spectra of noise in the quadrature amplitudes of light transmitted through the cavity. Starting from the bottom, the three subplots show

the spectrum recorded with a probe power that increases in 4 dB increments. The optical local oscillator power was 1 mW per photodetector, and the bottom trace shows the probe power level corresponding to $\langle m \rangle = 1.5$ photons in an empty cavity. One can clearly see the noise peaks carried by the probe beam rising up out of a flat background of shot noise, or at least noise carried by the optical local oscillator (see below). The most prominent features in the noise spectrum are the set of peaks between 10 and 14 kHz ($\sim 0.06 f_{Nyq}$), a large peak near 32 kHz ($\sim 0.16 f_{Nyq}$), and another large peak near 54 kHz ($\sim 0.27 f_{Nyq}$). One might ask whether these are contributed mainly by vibrations of the physics cavity, or by frequency noise on the Ti:Sapph. I think there is strong evidence for the former. First off, one sees conspicuous features at both 10 and 14 kHz in my initial measurement of the physics cavity transfer function. For the measured Ti:Sapph closed-loop noise spectrum, however, things are quiet in that range, and it would look as if the tweeter resonances at 17-20 kHz should really be the first things to show. Next, the same set of major peaks shows up in the noise spectrum of the locking-diode/physics cavity error signal (Figure 2.4). By taking some simultaneous recordings of the error signal and phase quadrature amplitude, I found that the Fourier signals at these frequencies are in fact phase coherent. This actually implies that one can subtract a bandpass-filtered version of the simultaneously recorded error signal from the phase quadrature to suppress the 10 and 14 kHz noise peaks. I found that a two-fold reduction of the power in these peaks can be accomplished, but even this turns out to have a small enough effect on the overall noise variance that it's not really worth the trouble.

Another benefit of having taken simultaneous recordings of the heterodyne photocurrent and physics-cavity error signal was the opportunity to verify directly that the atomic transits don't effect the physics cavity servo.

2.2.4 Cesium MOT

To provide a source of cold Cs atoms, we used a standard magnetooptic trap loaded directly from a thermal beam [57]. Our choice of thermal-beam loading, as opposed

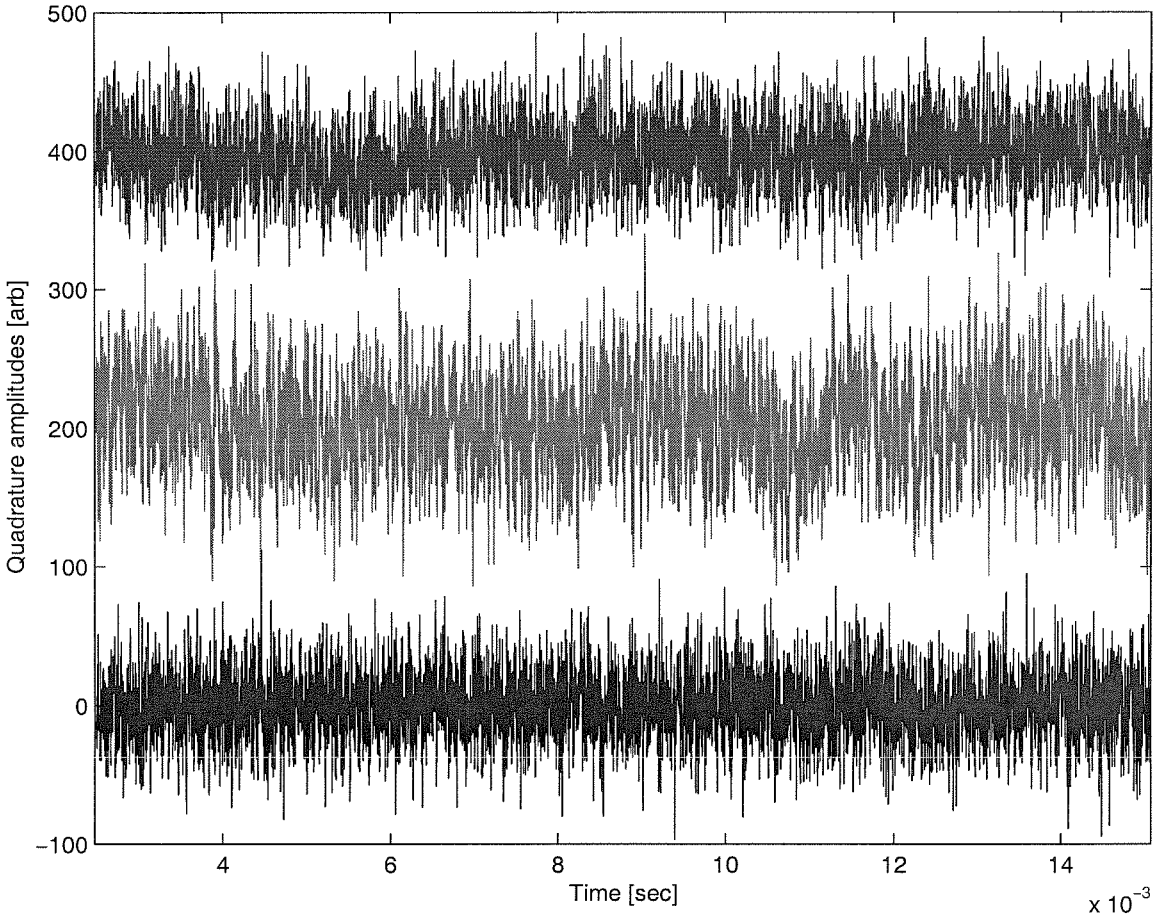

Figure 2.2: Examples of demodulated signals from the local oscillator alone (BOT), the phase quadrature of a transmitted probe beam (MID), and the amplitude quadrature of a transmitted probe beam (TOP). The latter two traces were recorded simultaneously.

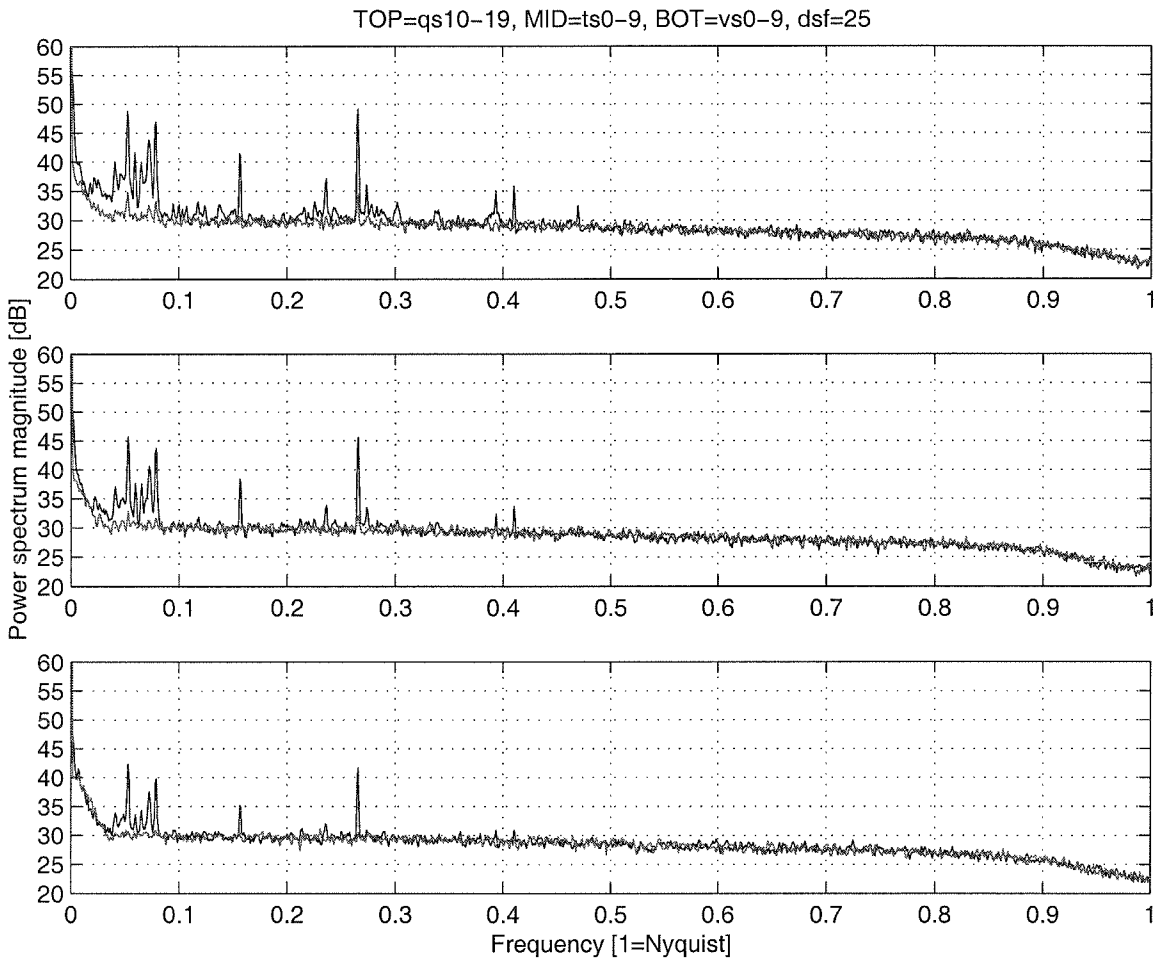

Figure 2.3: Power spectra of noise on amplitude and phase quadratures (see text). In each subplot the, the trace with noise peaks is the phase-quadrature noise spectrum. The Nyquist frequency is 200 kHz.

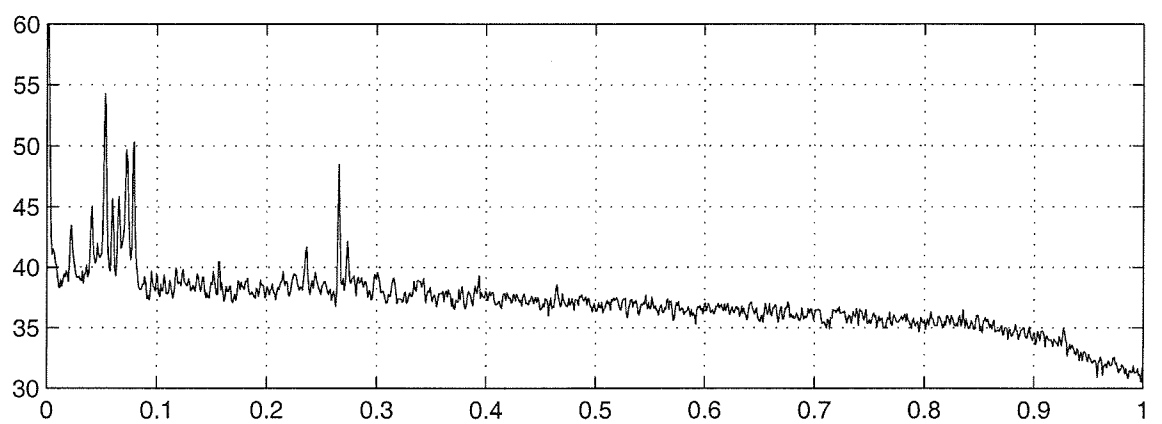

Figure 2.4: Power spectrum of noise on physics cavity error signal.

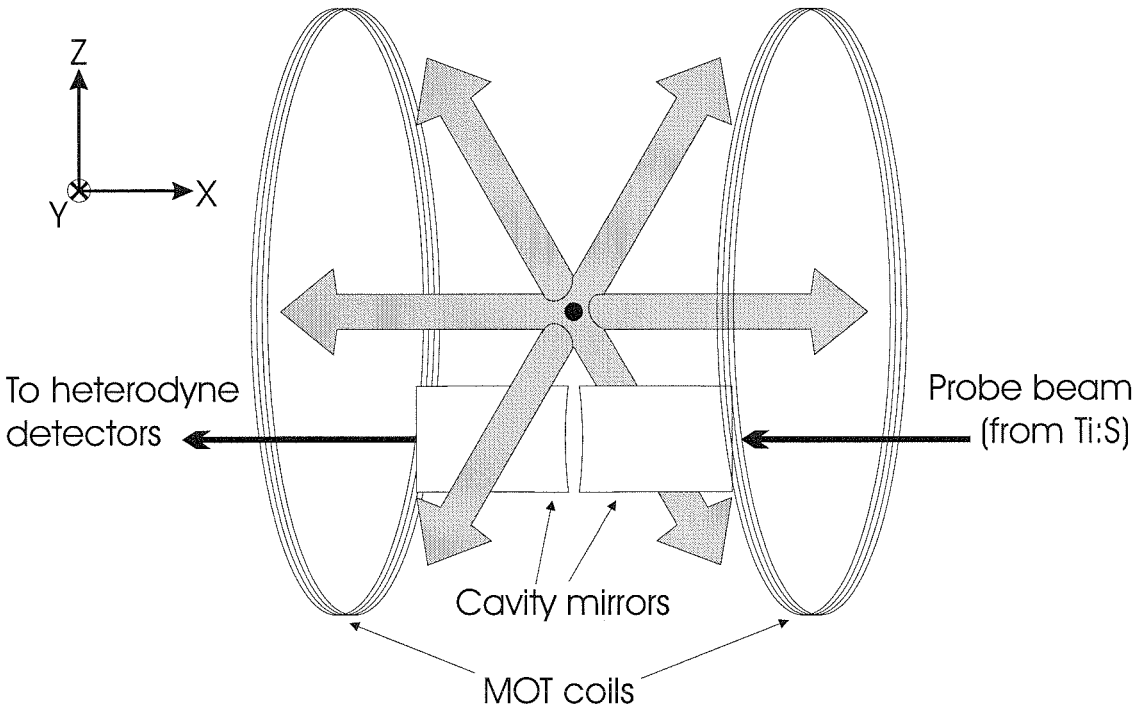

Figure 2.5: Geometrical arrangement of the MOT beams, relative to the physics cavity (drawing not to scale). The mirror substrates are each 3 mm in diameter and 4 mm long. The MOT forms at a height of ~ 7 mm above the cavity axis.

to loading from a background vapor, was driven by an attempt to prevent accidental coating of the physics cavity mirrors with Cs. In more than two years of service, we did not detect any significant ($> 5\%$) change in the cavity finesse. Pre-cooling of the Cs beam was not necessary for this experiment, as we required only a very low rate of delivering single cold atoms into the cavity mode volume. Our MOT employs a six-beam configuration, and we orient the anti-Helmholtz coils for the trap so that their symmetry axis is parallel to that of the optical cavity. This leads to a MOT laser beam geometry with one beam axis running parallel to and just above the cavity, plus two beam axes in the plane of the mirror surfaces (Figure 2.5).

The light for the MOT was provided by a pair of grating-stabilized diode lasers (SDL 5421-G2), one tuned to the Cs $6S_{1/2}(F = 4) \rightarrow 6P_{3/2}(F' = 5)$ cycling transition

for trapping and the other to $6S_{1/2}(F = 3) \rightarrow 6P_{3/2}(F' = 4)$ for repumping. The trapping beams were $\sim$ 1 cm diameter with anywhere from 40 μW to 4 mW of power, depending on how many atoms we were trying to send into the physics cavity. We typically used a Cs reservoir temperature of 60-80 C for the thermal beam, which effused through a 200 μ pinhole and travelled an overall distance of $\sim$ 60 cm to the trapping region (with a cold mechanical collimator in the way to reduce loading of the ion pump). With an anti-Helmholtz field of around 25 G/cm, we could load up to $\sim 2 \times 10^5$ atoms into a millimeter-sized cloud, whose mean temperature we estimate to be $\sim$ 100 μK based on fluorescent imaging of free expansion (Figure 2.6) and on the spread in arrival times of individual atoms falling into the cavity (Figure 1.3). For the MOT-expansion measurement, snapshots of the atom cloud were taken at successive time delays using a gated intensified CCD camera and PC-based frame grabber. An acousto-optic modulator was used to produce brief ($\sim 1\mu$s) pulses of diode laser light to excite the atoms, with the $\sim$ 100 ns wide gate for the CCD intensifier synchronized to the leading edge of the light pulse.

When running the experiment we would load the MOT for about 0.5 s, then drop it by quickly turning off the trapping beams with an AOM (using an rf switch with $\sim$ 45 dB attenuation). After the trapping beams are thus extinguished, we ramp down the anti-Helmholtz field according to an RC-filtered step with $\sim$ 4 ms time constant. The repumping beam was left on all the time, so that falling atoms would be shelved in the $F = 4$ ground hyperfine level before entering the cavity. No specific preparation was performed with respect to the atomic Zeeman states. Dropping $\sim 2 \times 10^5$ atoms we would generally see 30-50 atoms falling through the central part of the cavity mode volume, so for single-atom transit data we had to reduce the trapping beam power by a factor of 100 to reach 0-2 atoms per drop. The overall repetition rate for the trap-drop cycle is typically 0.6 Hz.

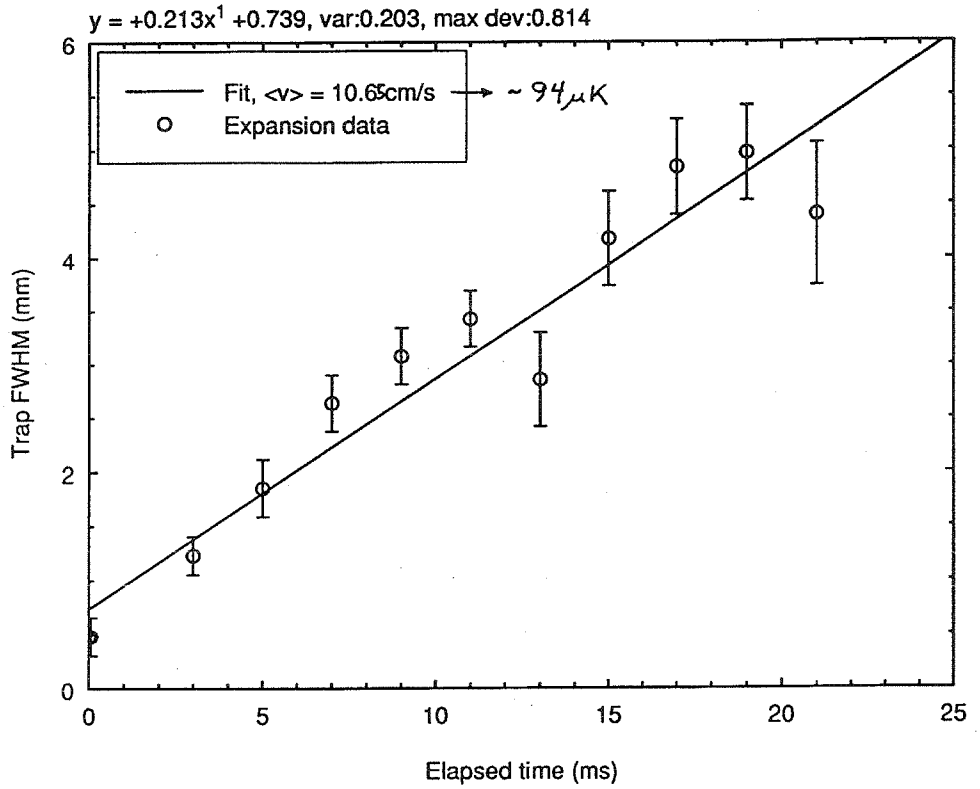

Figure 2.6: Estimate of initial MOT temperature based on rate of free expansion.

2.2.5 Probe generation and photodetection

We use a balanced super-heterodyne setup in order to achieve high-efficiency, zero-background photodetection of $\sim$ pW levels of 852 nm light transmitted through the physics cavity. The frequency difference between cavity probe light and the optical local oscillator for heterodyne detection was between 40-190 MHz, depending upon our choice for the atom-probe detuning. The probe light was generated from the Ti:Sapphire output by cascading a +200 MHz AOM and a tunable travelling-wave electrooptic modulator, which was driven between -245 MHz and -440 MHz to produce the desired atom-probe detuning. This indirect method was required to prevent contamination of the photocurrent by electronic noise at the heterodyne frequency.

Light leaving the physics cavity first hits a color-separation mirror which reflects $> 99\%$ of the 852 nm light but transmits $\approx 30\%$ of the 836 nm light, allowing us to recover an error signal for locking the physics cavity (see above) without compromising the overall detection efficiency for the probe field. Residual 836 nm light going to

the heterodyne setup amounts to only $\sim$ 30 nW and has a negligible effect on the photocurrent of interest.

The local oscillator (LO) for the optical heterodyne is spatially cleaned by a $\mathcal{F} \sim$ 1000 Fabry-Perot cavity (linewidth $\sim$ 1 MHz), which also serves to strip off spectral noise at 76 MHz associated with the AOM servo for stabilization of the Ti:Sapphire frequency. The cleaning cavity is locked using the Pound-Drever-Hall method [54] with FM sidebands at 24 MHz, which likewise must be kept weak in order not to saturate the AC gain of the heterodyne photodetectors. We use a total of $\sim$ 2 mW in the LO, which generates a shot-noise level $\sim$ 5 dB above the electronic noise of the photodetectors in the frequency range of interest.

Our photodetectors incorporate Hammamatsu S5972 Silicon PIN-diodes (Q.E. $\approx$ 70%), and a Signetics NE5211 transimpedance amplifier as the first rf gain stage. While this design achieved high sensitivity over nearly 300 MHz bandwidth, great care had to be taken to surpress self-oscillations. It appears that the additional circuitry required by such complications led to slight differences in the roll-off of one detector relative to the other, making broadband balancing of the heterodyne rather difficult. We can achieve $\sim$ 20 dB balancing in narrow bandwidths, however, as measured by the cancellation of sinusoidal amplitude modulation applied directly to the local oscillator. In any case, LO technical noise at the heterodyne frequency was not an important issue because the frequency was so high—the main importance of balancing was to prevent loss of heterodyne efficiency due to amplitude/phase errors in the subtraction of the two photocurrents.

The difference photocurrent from the balanced heterodyne detectors was amplified up to -50 dBm or higher, then divided by a 0° rf splitter. An independent signal generator was used to produce an rf local oscillator at the heterodyne frequency, and it was halved using a 90° rf splitter. The two identical copies of the photocurrent were mixed with the in-phase and quadrature copies of the rf LO to produce an orthogonal pair of quadrature amplitude (QA) signals at baseband. The QA signals were further amplified, and passed through 300 kHz analog filters with a roll-off of 12 dB/octave. We used a 12-bit ADC to sample both QA's simultaneously at a rate of 10 MHz per

channel, which is sufficiently high to avoid signal aliasing completely. Following each drop of the trap, we continuously recorded both QA's for a data acquisition window of 50 ms and streamed the data to a hard drive for offline processing following the experimental run.

Ideally, we would like the data acquisition procedure just described to yield directly the amplitude and phase quadrature amplitudes of light transmitted through the cavity. If we write the transmitted optical field as $\mathcal{A}(t)e^{-i\omega_p t}$, where $\mathcal{A}(t)$ is a slowly-varying complex amplitude, the amplitude x_a and phase x_p quadrature-amplitudes are defined by $\mathcal{A} \equiv x_a + ix_p$. With respect to the quantum-mechanical theory of the Master Equation (2.3), $\mathcal{A} \propto \langle \hat{a} \rangle$. Note that I am defining $\mathcal{A}$ to have zero phase when the cavity is empty, so that x_p should have zero mean when there are no intracavity atoms.

Given the way that we generate the probe beam, however, we have no way of generating a phase-locked rf local oscillator to recover x_a and x_p directly. The phase of the heterodyne photocurrent differs from the phase that the light has just after it leaves the cavity because of fluctuations in the optical path length from the cavity to the photodetectors. So the in-phase and quadrature signals produced by mixing the photocurrent with our rf local oscillator correspond to an orthognal, but rotated pair of quadrature amplitudes x_1, x_2:

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} cos\phi & -\sin\phi \\ \sin\phi & \cos\phi \end{pmatrix} \begin{pmatrix} x_a \\ x_p \end{pmatrix}. \tag{2.14}$$

Luckily, the characteristic timescales for fluctuations in the phase ϕ are quite long ($\sim 10 - 100$ ms, corresponding to acoustic disturbances) compared to the 250 μs duration of an individual atom-transit signal. In processing the recorded data to produce the plots discussed below, I have therefore used an "adaptive" definition of the amplitude and phase quadrature amplitudes. Within a window of 2 ms preceding the signal of interest, I estimate the instantaneous value of ϕ by determining the rotation of x_1,x_2 that produces one quadrature $\tilde{x}_p$ with zero mean and one quadrature $\tilde{x}_a$ with positive mean. Then $\tilde{x}_p$ is operationally defined to be the phase quadrature,

and $\tilde{x}_a$ is the amplitude quadrature.

2.3 Determination of intracavity photon number

In order to make quantitative comparisons between data and theory, we need to calibrate the strength of the driving field used in each experimental run. As we will ultimately choose to ignore DC optical phases, the relevant quantity for us will be the number of photons $\langle m \rangle$ that builds up when the cavity is empty. This corresponds to $2\kappa_a |\mathcal{E}|^2 / (\kappa_a + \kappa_b + \kappa_c)^2$ in the Master Equation case and $|y|^2/m_0$ for the OBSE.

Our strategy for determining the intracavity photon number during experimental runs has been to work backwards from the heterodyne photocurrent observed with no atoms in the cavity, using the expression [46]

$$\left(\frac{S^2}{N}\right)_{het} = 2T\eta \left(2\kappa_b \langle m \rangle\right), \tag{2.15}$$

where S is the peak amplitude of a demodulated quadrature amplitude signal (*i.e.* the heterodyne photocurrent looks like $S\cos(\omega t + \phi)$, N is the mean-squared power of fluctuations in the same signal due to (optical) shot noise, and T is the measurement interval. Note that this expression is valid for a coherent signal beam, which we assume to be the state of the light transmitted through the cavity when no atoms are present. For an accurate calibration, we thus need to know the output mirror transmission κ_b, the total cavity loss (which can be determined from measurements of the cavity finesse), and the overall photodetection efficiency η. In determining $\langle m \rangle$ from the data, we typically chose T on the order of 1 ms. The following subsections provide further detail for each critical aspect of the calibration.

2.3.1 Evaluation of local oscillator noise

Ideally, the quantity N appearing in equation (2.15) should coincide with the photocurrent noise power n observed when the signal beam is blocked. This should allow us to determine the intracavity photon number without having to calibrate the exact

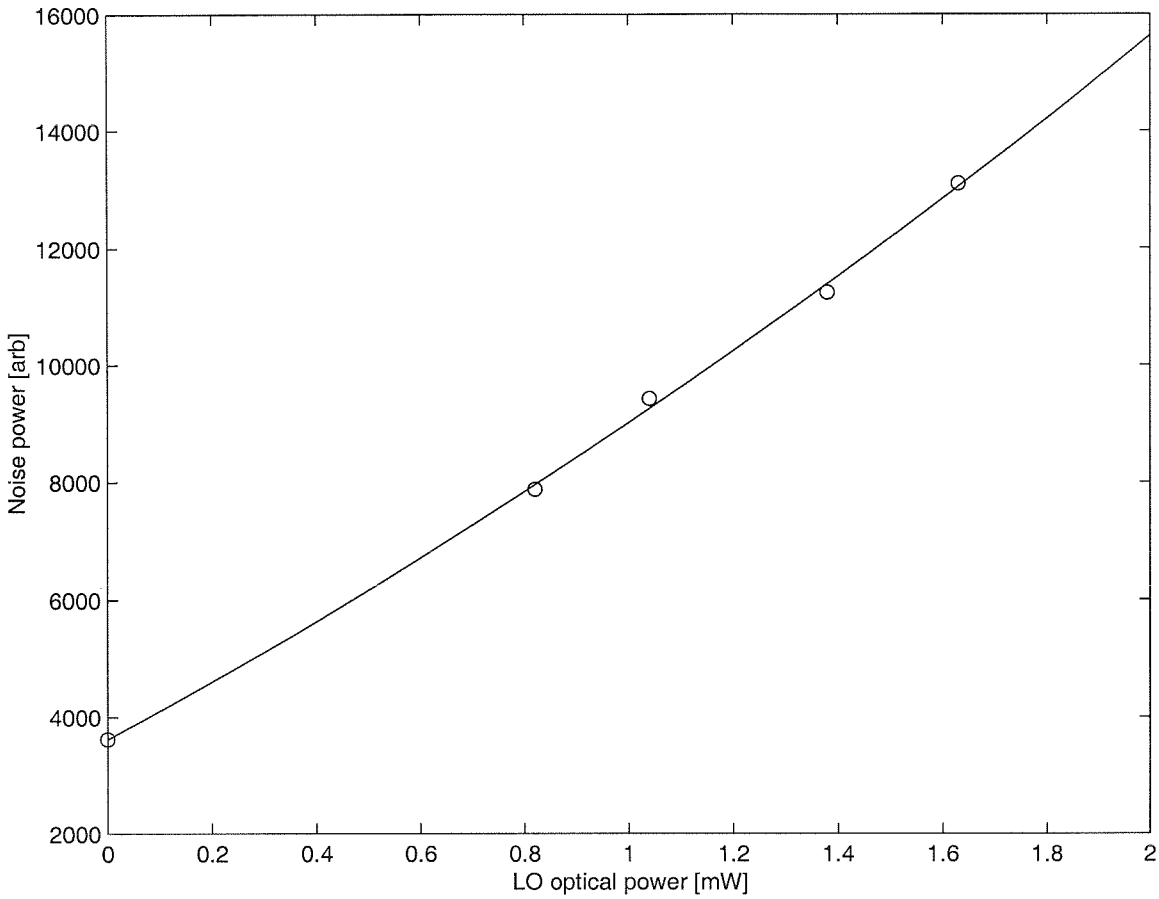

Figure 2.7: Plot of the photocurrent noise power versus LO optical power, with no signal beam to the heterodyne detectors.

gains of the photocurrent amplifiers, etc. But the procedure is invalid if the optical local oscillator carries excessive technical noise, so we briefly examined the scaling of our nominal N with the DC optical power of the LO [58]. Figure 2.7 shows several data points of the photocurrent noise power n (300 kHz analog bandwidth) versus LO power P. A fit of the data to the functional form $n = aP + bP^2$ yields $b/a \simeq 0.11$ mW^{-1}. Given our typical operating LO power of 1 mW per detector, and considering the relative magnitudes of other uncertainties, we approximate $N \simeq n$. Note that the LO does pass through an intensity stabilizer with $\sim$ 400 Hz bandwidth just before entering the cleaning cavity.

2.3.2 Measurement of cavity decay rates

We directly determined the total cavity field decay rate $\kappa_a + \kappa_b + \kappa_c$ by making a calibrated measurement of the HWHM of a cavity TEM_{00} resonance. In order to minimize systematic errors, we did this by using the heterodyne detectors to monitor the transmitted optical power of a fixed-frequency Ti:Sapph probe beam while scanning the cavity length. The cavity length was always under servocontrol during the measurement, as we generated the scan by stepping the frequency of the rf going to the travelling wave modulator for the locking diode laser. The total cavity field decay rate is then given by the resonance HWHM measured in terms of the modulator rf scan, times a correction factor of the ratio of the diode laser and Ti:Sapph wavelengths. We find $\sum \kappa \simeq 3.2$ MHz.

As the mirrors used to construct the cavity should be identical, we assign half the total losses to each mirror. Unfortunately, we did not manage to characterize the ratio of intracavity losses κ_c to transmission losses $\kappa_a + \kappa_b$ before we accidentally damaged the cavity. Given that the mirror coatings we have previously received from the same manufacturer have displayed very low scattering/absorption loss in the wavelength range of interest [59], we have assumed $\kappa_b = 1.6$ MHz.

2.3.3 Measurement of η

Three principle factors determine the value of $\eta \equiv \mathcal{V}\varepsilon\tau$: the spatial overlap $\sqrt{\mathcal{V}}$ between the signal beam and the optical local oscillator, the photodetector quantum efficiency ε, and the (power) efficiency τ with which we transfer light from the cavity ouput to the photodetectors. Using an optical powermeter we measured $\tau = 0.9$, as well as $\varepsilon \simeq (0.68, 0.71)$, for our two photodetectors. We measured $\mathcal{V} \simeq 0.65$ by producing a DC fringe between the signal beam and the optical local oscillator, with both beams adjusted to a power level of 1 μW. The power adjustments were made upstream of the physics cavity and LO cleaning cavity, ensuring that no optical misalignments were introduced in the process. We obtained an independent measurement of $\mathcal{V}\varepsilon \simeq 0.4 \pm 0.05$ (where the quoted uncertainty indicates repeatability)

from the heterodyne signal-to-noise ratio, as measured by our HP Spectrum Analyzer [18], obtained with a nominal signal power of 10 nW (measured with an optical power-meter). Note that the reasonable agreement between this latter value and the directly-measured $\mathcal{V}\varepsilon$ provides further indication that our optical local oscillator bears minimal excess noise.

Although it could have been avoided, we did suffer one additional loss of detection efficiency due to gain and phase imbalance in the subtraction of heterodyne photocurrents. Several weeks after we took the data sets shown below, we realized that one photodetector output had an rf signal 4 dB higher than the other (this figure includes both the deviation of our heterodyne beamsplitter from being 50/50, the difference in photodiode quantum efficiencies, and the difference in transimpedance gains), and a phase offset of 1 radian. If we write the two photocurrents (for a coherent signal beam) as $i_1 = \frac{1}{\sqrt{2}}e^{-i\omega t} + \xi_1$ and $i_2 = \frac{-g}{\sqrt{2}}e^{-i\omega t + i\phi} + g\xi_2$ (where ξ_1 and ξ_2 are uncorrelated gaussian white noises with zero mean and unit variance), the difference photocurrent is

$$i_- = \frac{1 + ge^{i\phi}}{\sqrt{2}}e^{-i\omega t} + \sqrt{1+g^2}\xi, \tag{2.16}$$

where ξ is again a gaussian white noise process with zero mean and unit variance. The complex imbalance $ge^{i\phi}$ should thus reduce the effective photodetection efficiency (for the purpose of evaluating the heterodyne signal-to-noise ratio (2.15)) by

$$\eta \rightarrow \frac{1}{2}\eta\frac{|1 + ge^{i\phi}|^2}{1+g^2}. \tag{2.17}$$

Using the measured $g \simeq 0.63$ and $\phi \simeq 0.85$ rad, $\eta \rightarrow 0.8\eta$, and we quote an overall value of $\eta \simeq 0.32$.

2.4 Numerical simulations

In order to facilitate the interpretation of our data, I wrote and ran a set of rudimentary Monte Carlo routines to simulate the heterodyne signals that we should see as atoms fall through the cavity. The codes simulate three-dimensional, classical

center-of-mass motion of individual atoms under the influence of forces due to their strong interactions with the cavity field. Gravity is also included in the kinematics, but plays only a minor role over the ~ 1 ms duration of the simulations.

2.4.1 Overall scheme for the simulations

The basic scheme of the simulations is to precompute the values of the mean cavity-induced force, the cavity-field-induced diffusion coefficient, and the steady-state (complex) amplitude for the intracavity field as a function of atomic position. Hence, *every* other degree of freedom in the simulation is adiabatically eliminated and slaved to the atomic motion, which is assumed to be the slowest and the "stiffest" process in the dynamics. The mean force and intracavity field are derived from steady-state solutions of the Master Equation (2.3), and the cavity diffusion is computed using the Quantum Regression Theorem [9]. In each timestep, the code first performs an interpolation on the precomputed tables of values to determine the appropriate change in atomic momentum (which includes a stochastic increment consistent with the local value of the diffusion constant), records an appropriate value for the cavity output field, and then updates the atomic position and momentum. The simplest possible integration scheme is used (explicit Ito-Euler), yielding order 0.5 convergence in the timestep. I used a timestep of 7.5 ns in order to keep the run-times for the simulations reasonable under Matlab on our Pentium II workstations, and this should have been sufficiently small to keep the integration error below other sources of inaccuracy.

The principle shortcoming of this scheme is that it assumes the atomic velocity will remain small enough that variations in the coupling strength g will be negligible over timescales on the order of κ^{-1} and γ^{-1}. If this condition is violated, then the steady-state values of the mean cavity force and cavity output are no longer appropriate. Another major approximation was made in treating only a two-level atom—optical pumping and/or the associated opto-mechanical effects (*e.g.*, laser cooling) could certainly play an important role in determining the shapes of our observed transit signals. A third, probably less significant approximation is that I treat the stochastic process

associated with recoil from spontaneous-emissions as being statistically independent from the one associated with dipole-force fluctuations. The possibility that interesting effects could arise due to quantization of the atomic motion is of course ignored as well.

Despite this substantial list of caveats, it's not so easy to tell the simulated signals apart from the experimental data. If we were to take this similarity at face value, it would become very tempting to focus on certain characteristic "events" that occur in both the numerical and experimental signals. In particular, signals of both origins display sharp, isolated dips and steps that would seem to reflect something unusual about the atomic motion. Looking at the simulations, such events are always associated with sudden changes in the component of atomic velocity along the cavity axis. We'll look at some supporting examples in a moment, but let me first describe the numerical method for computing the cavity diffusion coefficients.

2.4.2 Computation of diffusion coefficients

Following Doherty *et al.* [29], I computed the diffusion coefficient associated with dipole-force fluctuations according to

$$D = \lim_{t\to\infty} \mathrm{Re} \int_0^\infty \langle \mathbf{f}(t), \mathbf{f}(t-\tau)\rangle d\tau, \tag{2.18}$$

where $\mathbf{f}$ is the force operator

$$\mathbf{f} = -i\hbar\nabla g\left(\mathbf{r}\right)\left(\hat{a}^\dagger\hat{\sigma}_- - \hat{a}\hat{\sigma}_+\right). \tag{2.19}$$

Note that the vector nature of the force operator comes only from the gradient $\nabla g\left(\mathbf{r}\right)$. Unlike the computation of mean forces and the expected intracavity field amplitude, the evaluation of (2.18) requires an actual time-integration of the cavity-QED Master Equation (2.3). Equation (2.18) is formally written as an expression in some sort of Liouville-Heisenberg picture, in that the force operator is assigned a time dependence

even though we're working with a dissipative system. What this really means is

$$\langle \mathbf{f}(t), \mathbf{f}(t-\tau) \rangle \equiv \mathrm{Tr}\left[\mathbf{f} e^{\mathcal{L}\tau} \mathbf{f} \rho_{ss}\right] - \langle \mathbf{f} \rangle^2, \tag{2.20}$$

where $\mathcal{L}$ is the Liouville superoperator defined by

$$\frac{d}{dt}\rho = \mathcal{L}\rho, \tag{2.21}$$

with reference to the cavity-QED Master Equation (2.3). Knowing the steady-state atom-cavity density matrix ρ_{ss}, one can evaluate the correlation function in (2.20) by integrating the Master Equation for a time τ with $\mathbf{f}\rho_{ss}$ as an initial condition, multiplying the result by ρ_{ss}, and finally taking the trace.

The numerical integration was done by a C++/MPI code that I wrote for SGI/Cray Origin-2000. I used a truncated basis of 25 Fock states for the cavity mode, using an explicit Euler integration with a 1 ps timestep. The total integration time is supposed to approximate ∞ in the upper limit of the integral in equation (2.18), and 5 μs seemed to be more than long enough. Running on sixteen nodes, each point in the lookup table took less than an hour to compute.

Figure 2.8 shows the results of such a calculation for the diffusion coefficient as a function of atomic position along the standing wave, for an atom located on the cavity axis. Recall that the cavity mode function varies as $\cos(k_L x)$ along the cavity axis, and as a gaussian in the transverse dimensions. Starting from the top, the four curves represent $\Delta = 0$, 10, 30, and 50 MHz, all with a probe strength such that an average of 2 photons would build up in the cavity if it were empty (and $\Theta = 0$). Figure 2.9 shows similar results, all for $\Delta = 50$ MHz, but with an average intracavity photon number (again in an empty cavity) of 8 for the top curve, 4 for the middle, and 2 for the lowest. As the variation of g along the standing wave is so much more rapid than in the transverse (gaussian) directions, I approximated $\nabla g(\mathbf{r}) \simeq dg/dx$ in expression (2.18). The computation of $\mathbf{f}$ for the mean force field was of course three-dimensional.

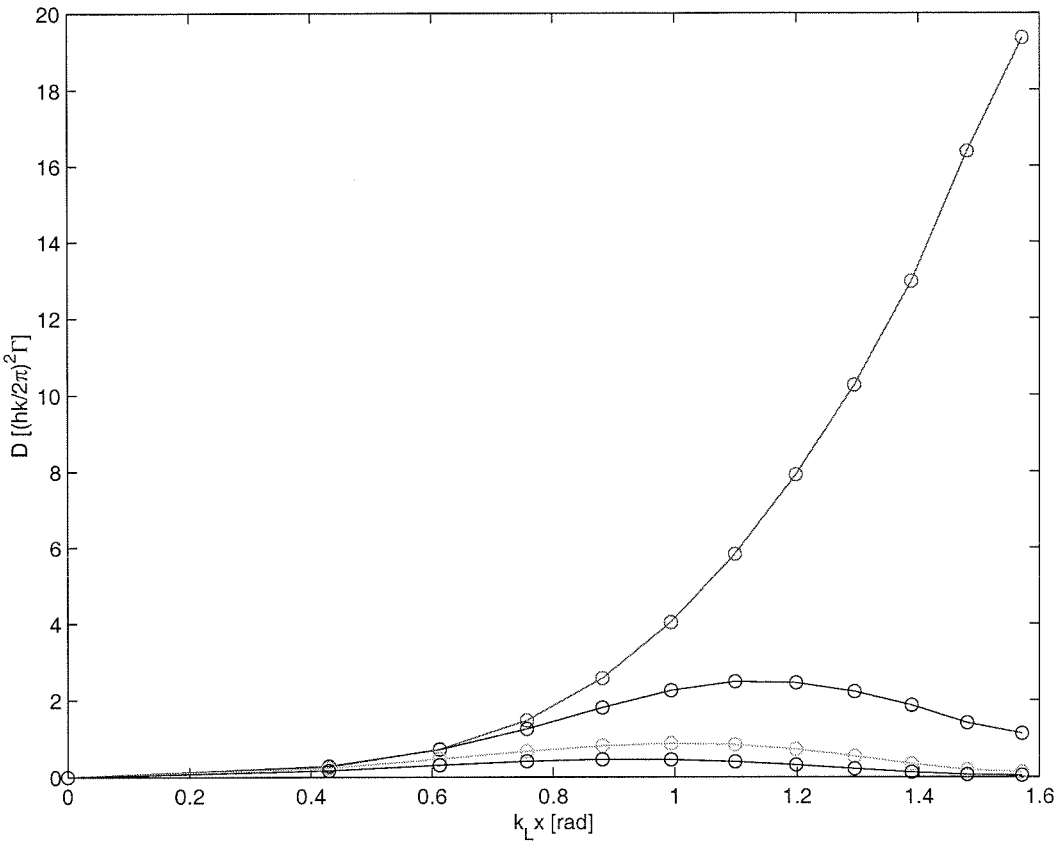

Figure 2.8: Cavity diffusion coefficients versus atomic position, for (from the top curve going down) atom-probe detuning $\Delta \equiv \nu_a - \nu_p = 0$, 10, 30, and 50 MHz, with $\langle m \rangle = 2$ photons.

Note that expression (2.18) gives the diffusion coefficient associated with dipole force fluctuations only—it takes no account of recoils from spontaneous emission. For the simulations, I therefore add a second diffusion process whose coefficient is computed from the expectation value of the atomic excitation at each point in space:

$$D_{rec} = \frac{(\hbar k_L)^2}{25} \Gamma \langle \hat{\sigma}_+ \hat{\sigma}_- \rangle. \tag{2.22}$$

Here Γ is the atomic spontaneous emission rate (Einstein A coefficient), and the number 1/25 comes from averaging over the angular distribution pattern for dipole radiation.

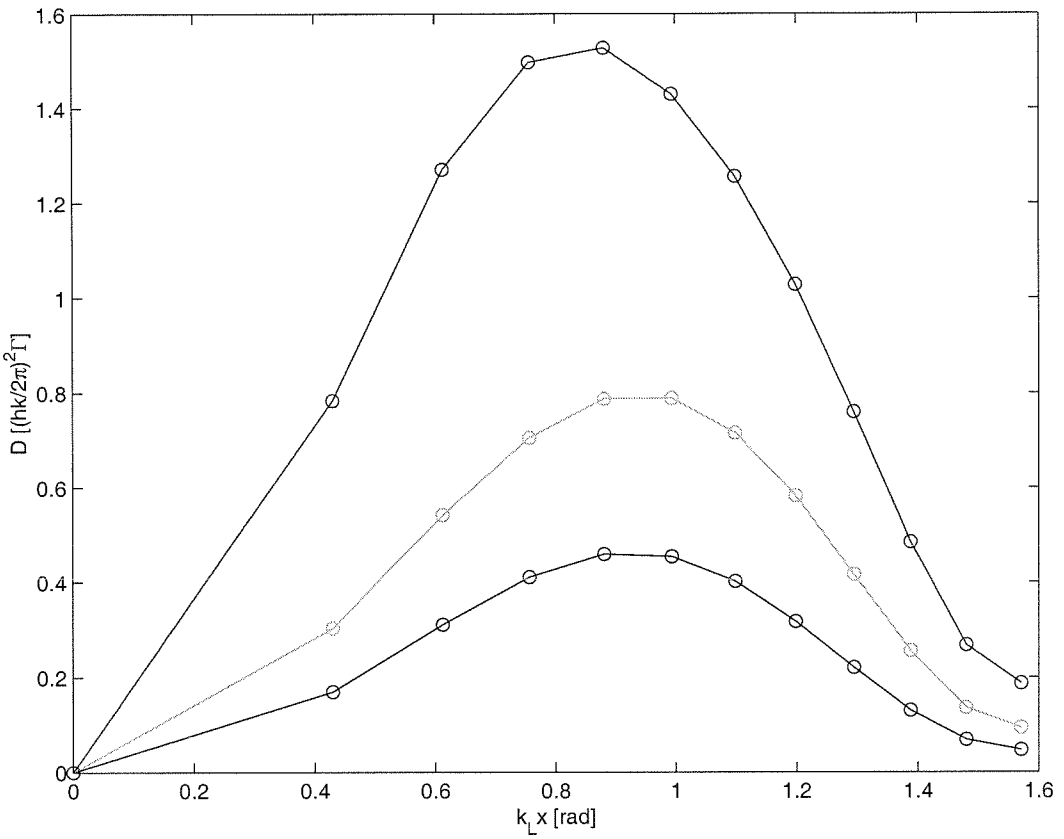

Figure 2.9: Cavity diffusion coefficients versus atomic position, for (from the top curve going down) $\langle m \rangle = 8$, 4, and 2 photons, with the atom-probe detuning $\Delta = 50$ MHz.

2.4.3 Discussion of results

Figure 2.10 shows a set of simulated signals for $\Delta = 10$ MHz and an average intracavity photon number of $\langle m \rangle = 2$. In each subplot, the upper trace shows the amplitude quadrature versus time, and the lower trace (with zero mean value) shows the phase quadrature versus time. Figure 2.11 shows similar signals for $\Delta = 50$ MHz and $\langle m \rangle = 4$. In all simulations, the atom was started at a position 7 gaussian waists above the cavity with a vertical velocity of -47 cm/s. The initial transverse position, transverse velocity, position along the cavity standing-wave axis, and velocity along the standing-wave were randomly assigned within parameter ranges that were narrow enough to ensure that most simulated atom-drops produced a sizeable signal.

The signals in Figures 2.10 and 2.11 were produced by taking records of $\langle \hat{a} \rangle$ versus t

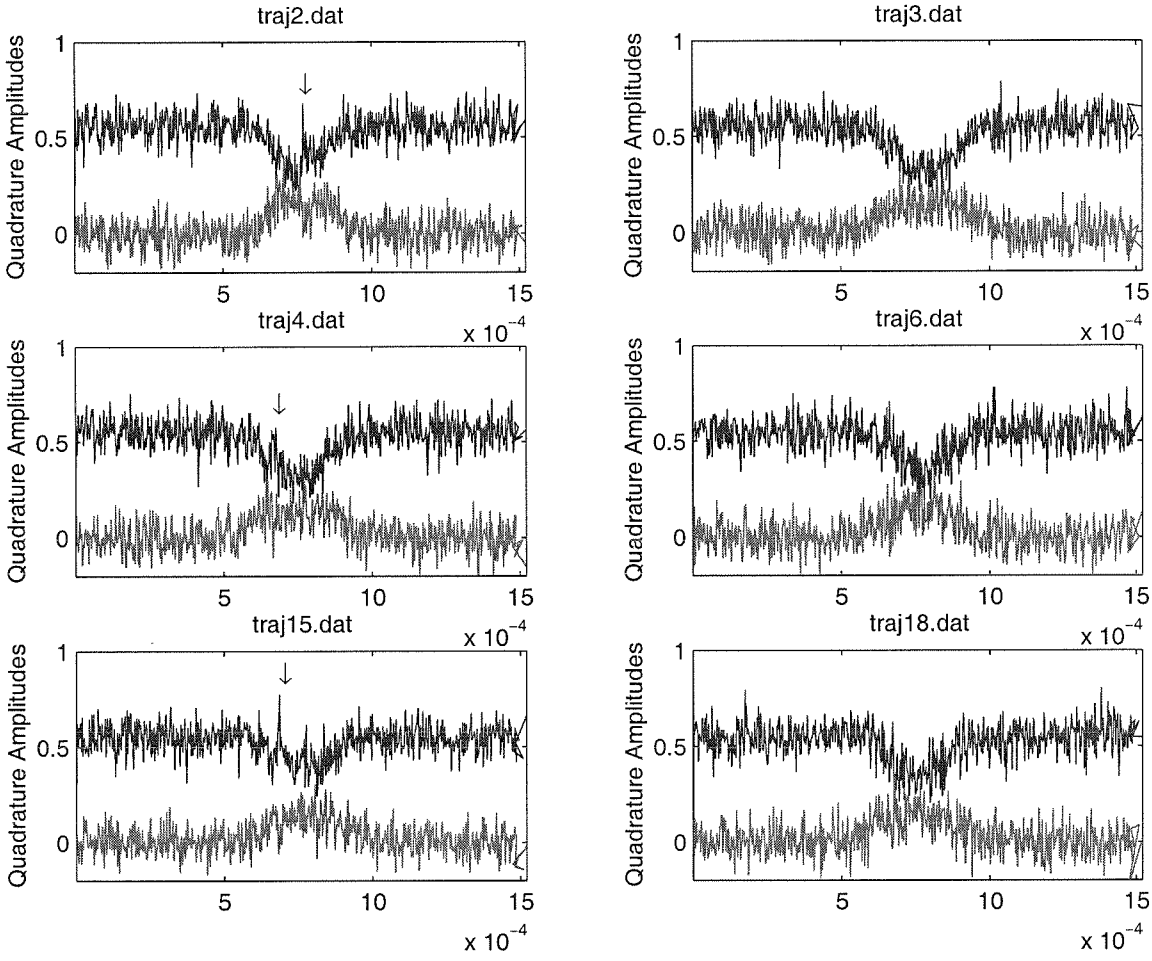

Figure 2.10: Some simulated atom-transit signals, for $\langle m \rangle = 2$ and $\Delta = 10$ MHz.

from the atomic motion simulator, adding shot noise to simulate photodetection with an overall efficiency of 0.32 as in the experiment (see equation (2.15)), and filtering down to a bandwidth of 200 kHz. They may directly be compared to the real data shown in Figure 2.18, although the range of the time-axes differ. One feature common to both the simulated and experimental signals is an occasional sharp dip, as in the plots labeled "traj2.dat," "traj4.dat" and "traj15.dat" (especially in the amplitude quadratures) at the points indicated by arrows. Similar dips can clearly be seen in the lowest right subplot of Figure 2.18 near 3×10^{-4} on the time axis, and also perhaps in the supblot just above it. Another recurring feature is the type of sudden step that occurs midway through the uppermost left subplot of Figure 2.11 (this time primarily in the phase quadrature), and also in the subplot just below it. With only a small

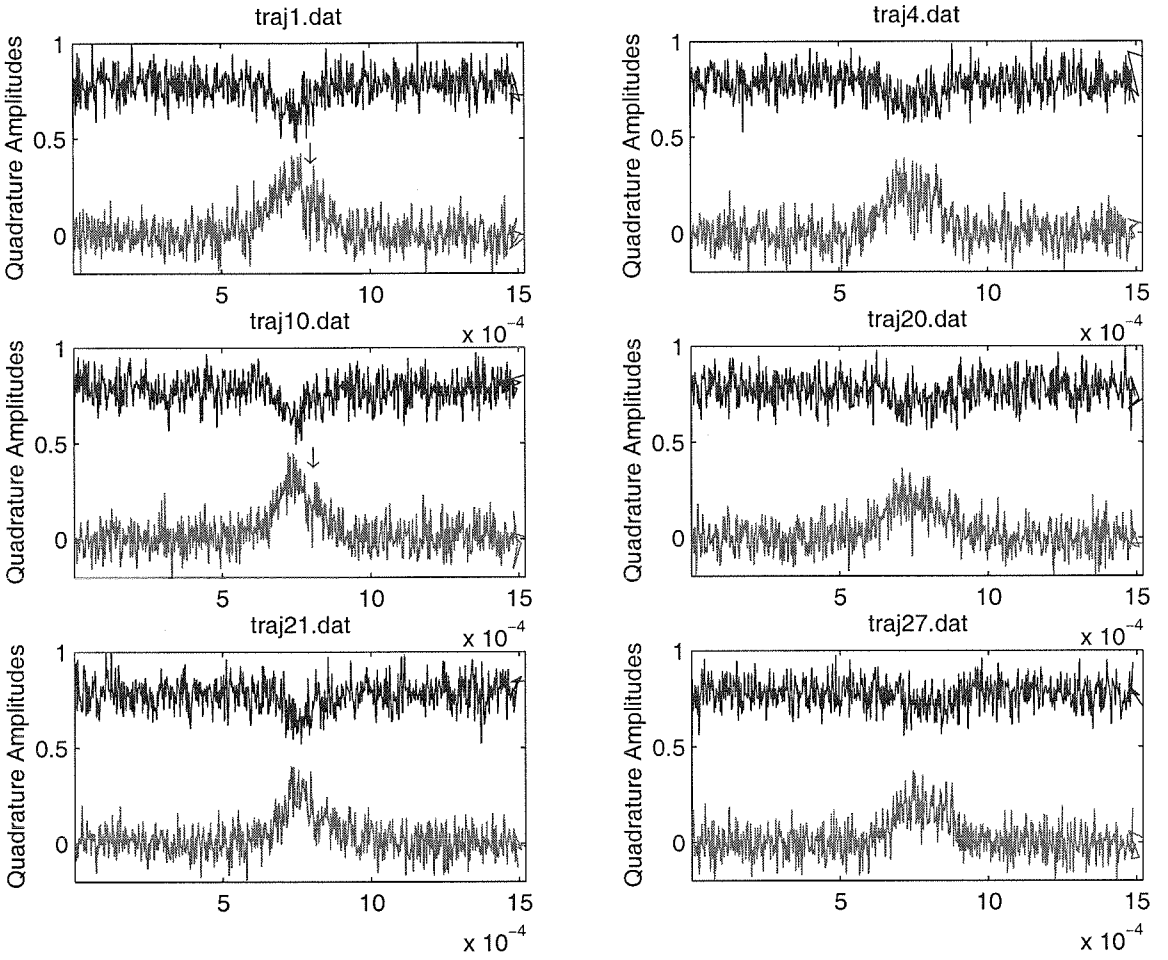

Figure 2.11: Some simulated atom-transit signals, for $\langle m \rangle = 4$ and $\Delta = 50$ MHz.

amount of imagination, similar steps can be seen in the two lowest subplots of Figure 2.18.

As mentioned above, it's difficult to know exactly what to make of such apparent similarities. One can look at the actual events in the atomic motion that caused the features to occur in the simulated data, but there's no way to know for sure whether similar events actually occurred in the experiment. In all fairness, however, it is rather difficult to come up with any explicit suggestions for optical pumping processes, *etc.*, that could provide an alternative explanation for the dips and steps. Although it's hard to quantify, such features really are endemic to the experimental data set.

Figures 2.12 through 2.15 show examples of the atomic position along the standing-wave x and the atomic velocity along the standing-wave v_x from four typical simu-

lations. Also shown are the mean value of the phase quadrature amplitude and a simulated signal with the appropriate amount of shot noise, both at a bandwidth of 300 kHz and of 150 kHz in order to indicate the effects of low-pass filtering on the qualitative features of the data. It certainly seems that the gross features of the simulated signals are dominated by atomic motion along the standing wave, in that the filtered versions tend to follow some "envelope" of it. The sharp dips would thus seem to be associated with moments at which the atom starts to accelerate along the standing wave, but does so slowly enough that the first few dips to $g = 0$ occur within the detection bandwidth. Sudden steps would correspond to more rapid accelerations such that the filtered signals go straight from full amplitude to half amplitude because of averaging over the standing-wave oscillations.

Looking through all the simulated data, it unfortunately seems that most of the virtual atoms are heated to sufficiently high velocities along the standing wave that the signals would oscillate at frequencies approaching 1 MHz. This implies velocities along the standing wave of 1 m/s or greater. Significant heating occurs even for the parameters with smallest cavity diffusion that I looked at in the simulations, $\Delta = 50$ MHz and $\langle m \rangle = 2$. For the highest bandwidth experimental data that we have, Fourier components in the range of 1 MHz would be attenuated by about a factor of ten and be lost in the rms noise. The oscillations don't seem to live in a sufficiently narrow frequency range that they could be pulled out by a bandpass filter, either. Of course, as discussed at the very beginning of this section, one has to doubt the real accuracy of the simulation strategy when the virtual atoms reach such high velocities along the standing wave. Neither does it seem that severe averaging over standing-wave oscillations is such a common problem in the experimental data, as a reasonable number of the observed signals seem to indicate atom-cavity couplings that reach the maximum value of $g = 11$ MHz with a nice, gradual envelope. Indeed, given the ~ 100 μm *length* of the cavity, real atoms in the experiment could only travel at 1 m/s speeds along the standing wave axis for ~ 100 μs before crashing into one of the mirrors. Figure 2.15 shows an example from the numerical simulations in which the virtual atom manages to remain at very low velocities along the standing wave.

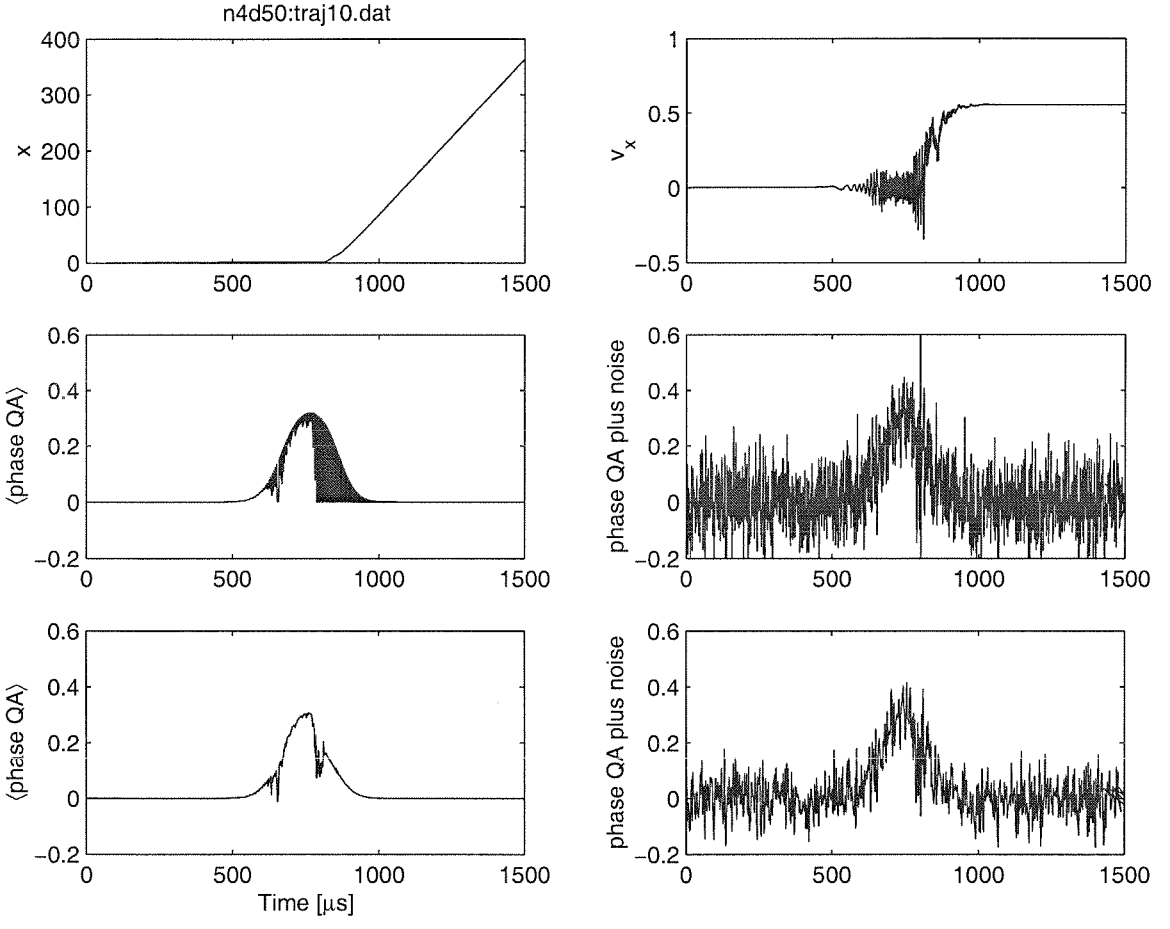

Figure 2.12: Simulated atomic trajectory and corresponding heterodyne signal for a single transit with $\langle m \rangle = 4$ and $\Delta = 50$ MHz. Here x is atomic position along the cavity standing-wave (measured in units of the optical wavelength) and v_x is the atomic velocity along the standing-wave (measured in optical wavelengths per μs).

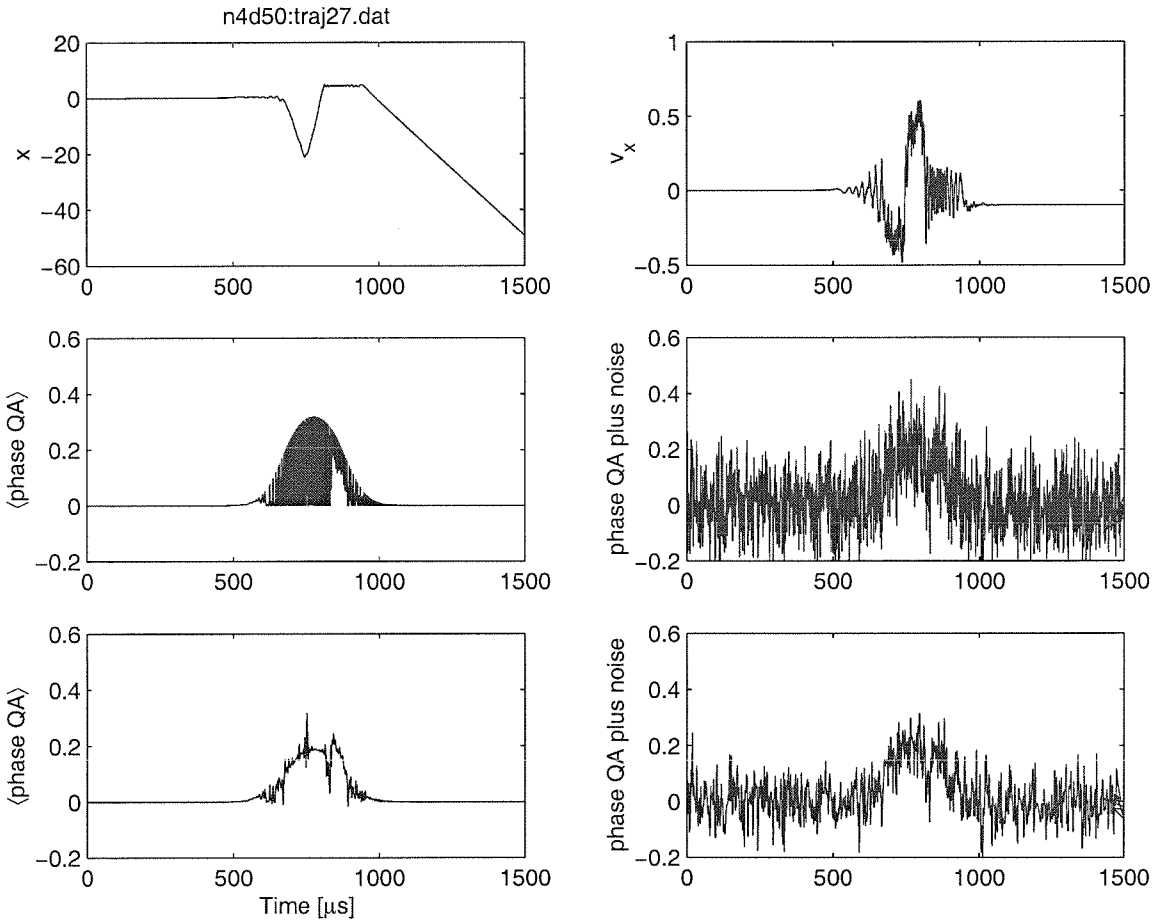

Figure 2.13: Simulated atomic trajectory and corresponding heterodyne signal for a single transit with $\langle m \rangle = 4$ and $\Delta = 50$ MHz. Here x is atomic position along the cavity standing-wave (measured in units of the optical wavelength) and v_x is the atomic velocity along the standing-wave (measured in optical wavelengths per μs).

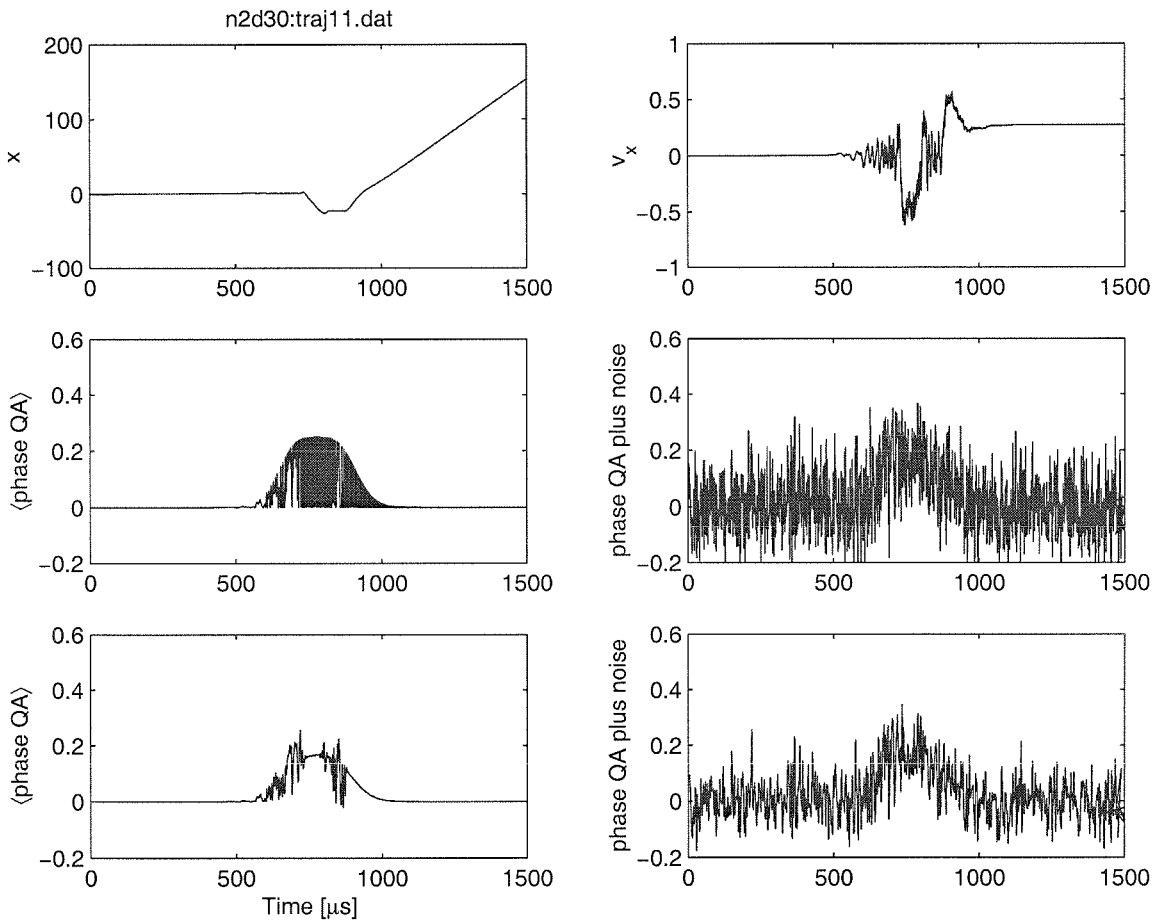

Figure 2.14: Simulated atomic trajectory and corresponding heterodyne signal for a single transit with $\langle m \rangle = 2$ and $\Delta = 30$ MHz. Here x is atomic position along the cavity standing-wave (measured in units of the optical wavelength) and v_x is the atomic velocity along the standing-wave (measured in optical wavelengths per μs).

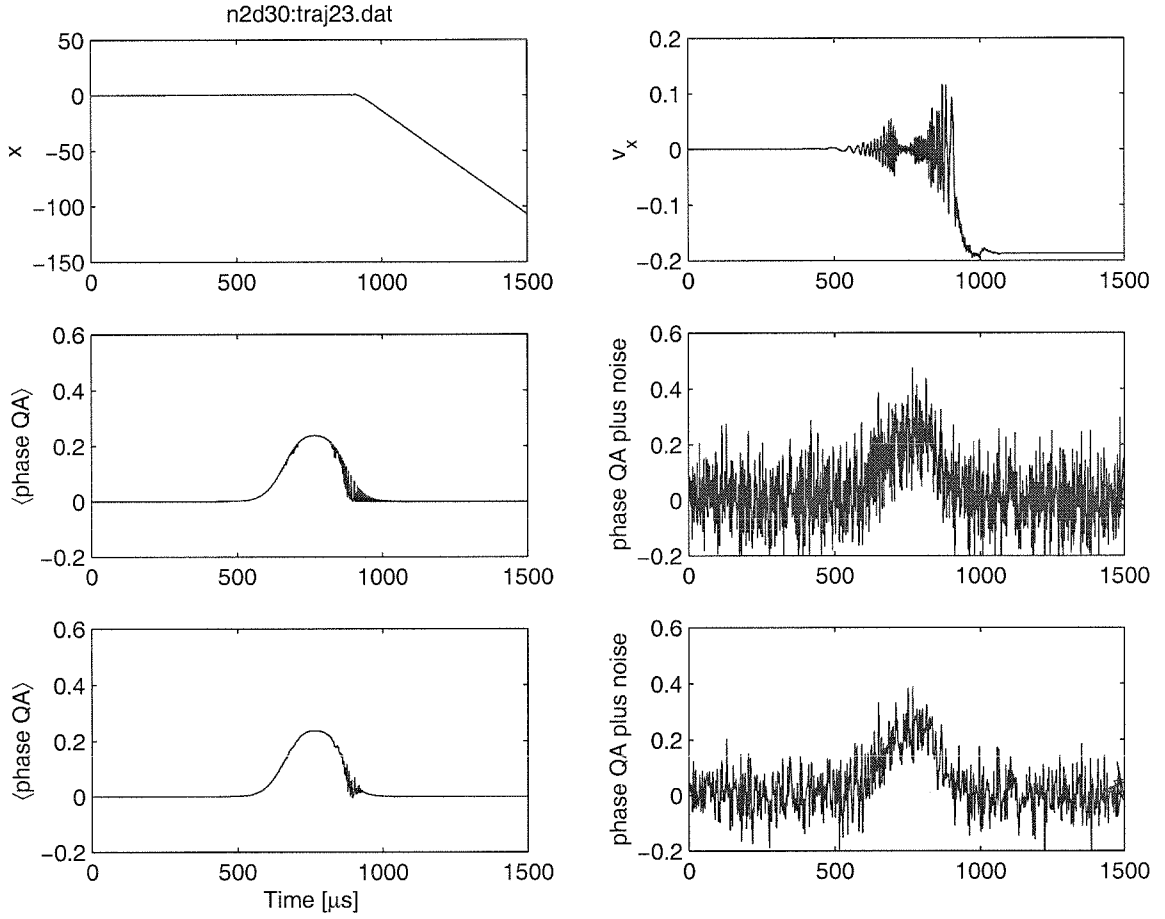

Figure 2.15: Simulated atomic trajectory and corresponding heterodyne signal for a single transit with $\langle m \rangle = 4$ and $\Delta = 50$ MHz. Here x is atomic position along the cavity standing-wave (measured in units of the optical wavelength) and v_x is the atomic velocity along the standing-wave (measured in optical wavelengths per μs).

2.5 Data

Moving on to the real data, I'll start by discussing some transit signals displayed in the same fashion as the simulations of Figures 2.10 and 2.11. The emphasis will be on trying to identify features in the experimental data that resemble the dips and steps we saw in the simulations. After that we'll look at some transit data displayed on the complex plane, for the primary purpose of making quantitative comparisons with theory.

Note that all of the data to be shown were taken with cavity-probe detuning $\Theta \equiv \nu_c - \nu_p = 0$, although the atom-probe detuning $\Delta \equiv \nu_a - \nu_p$ varies.

2.5.1 High bandwidth single-atom transits

The basic unit of our experimental data is a continuous stream of quadrature-amplitude values versus time, recorded for $\sim$ 50 ms windows following each dropping of the MOT. Several $\sim$ 15 ms segments of such data are shown in Figure 2.16. These were all taken with a detuning of $\Delta = 10$ MHz and with a probe power such that an average of 1.5 photons would build up in the cavity if it were empty (a situation I'll continue to denote $\langle m \rangle = 1.5$). One immediately notices some excess, low-frequency technical noise in the phase quadrature (the one with zero mean value) relative to the amplitude quadrature, but the correlation time of these fluctuations is sufficiently long that they don't really interfere with the atomic transit signals. The amplitude quadrature is quite clean, and as discussed above comes very close to the shot-noise limit for signal-to-noise ratio. I hope that the transit signals should be obvious—they are the narrow events clustered between $\sim 8 - 15$ ms on the time axis, in which the amplitude quadrature dips down and the phase quadrature comes up to meet it.

Zooming in on individual events, Figure 2.17 shows six of the largest signals from our data of December 22, 1997. The detuning and photon number for each signal are displayed by the y-axis. These are shown at our full analog bandwidth of 300 kHz (= -6dB point for power, implying that Fourier components of the signal at 300 kHz are attenuated by 0.5) and sampling rate of 10 MHz. The largest signals in the

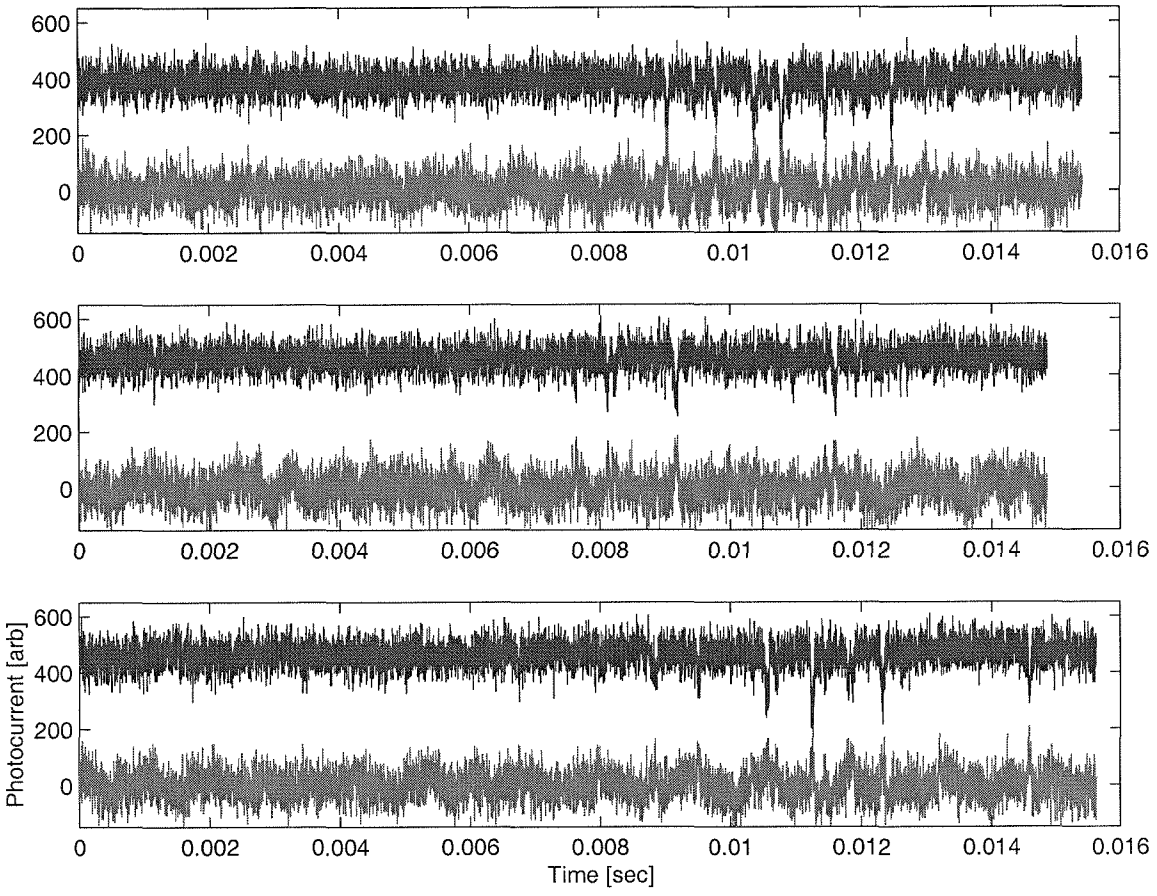

Figure 2.16: Three 15 ms segments of the recorded data.

data set tend not to have any significant internal structure of the type seen in the simulations, although one can see that the shapes of the overall signal envelopes do vary quite a bit. The distribution of the atom-transit signal between amplitude and phase quadratures clearly depends upon the probe detuning, as will be discussed in greater detail below. The lowest right subplot, taken with $\Delta = 50$ MHz and $\langle m \rangle \simeq 9$, actually represents an anomalously large event for the given parameters. Its size is slightly larger than the theoretical maximum for a single atom. We are therefore led to suspect that this event may have been caused by two atoms, or perhaps the major features of the signal were determined by one principle atom that happend to transit the cavity while several others were floating through the gaussian tails of the cavity mode. Generally speaking, we tried to keep the average atom number low enough so that this sort of thing would be minimized, but there's an obvious tradeoff between

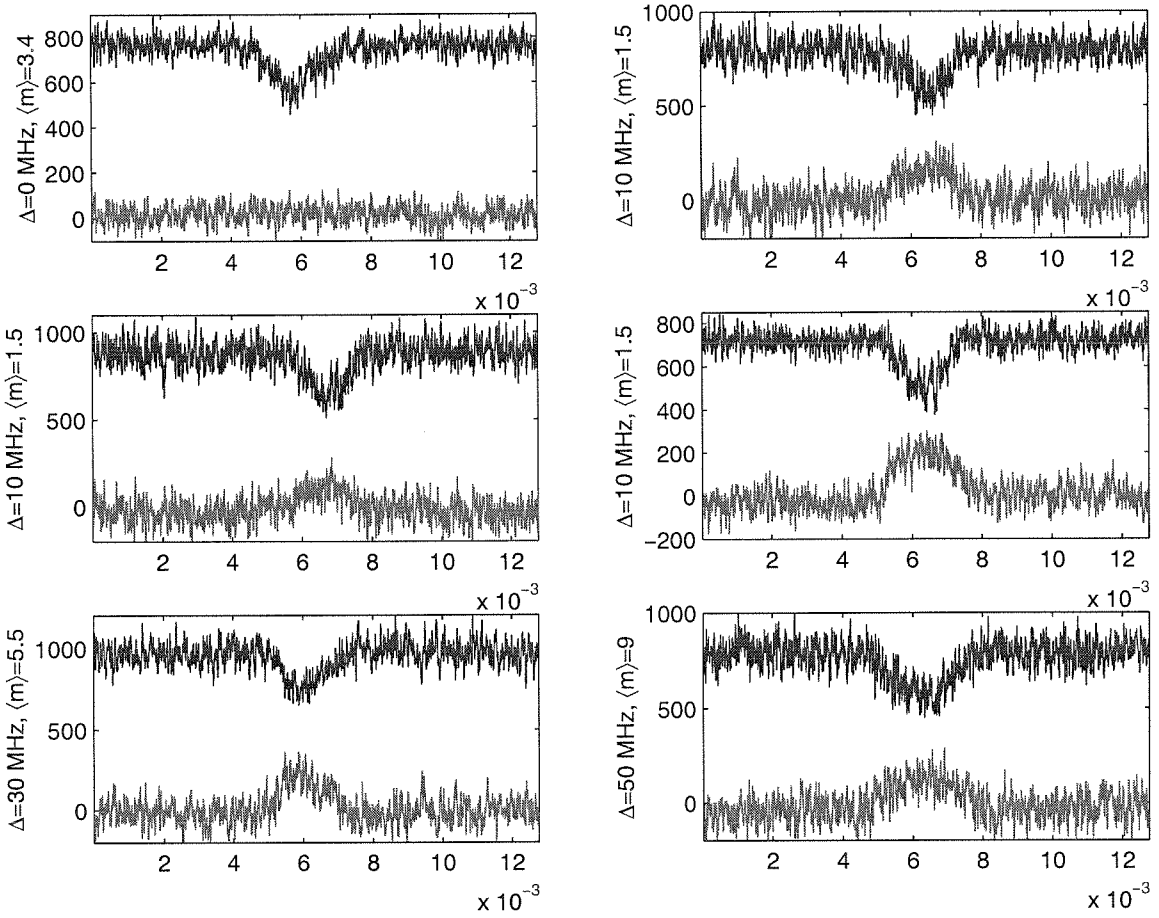

Figure 2.17: The largest atom-transit signals, shown at full bandwidth (see text).

obtaining isolated events and filling the hard drive with lots of empty traces. The data segments shown in Figure 2.16 reflect about the maximum average atom number per drop that we used in our experimental runs.

The only prominant feature to be seen in the high-bandwidth data of Figure 2.17 would seem to be the wiggles in the amplitude quadrature of the middle right subplot. This particular set of six transit signals was selected on the basis of signal "size," so it may be that these are really events in which the atoms maintained low velocities throughout their entire transits. For comparison, Figure 2.18 shows some signals that were selected specifically because they seemed to have some internal structure or an unusual overall envelope. Again, the detuning and photon number for each transit signal is displayed by the y-axis. As mentioned in the discussion of the simulated data, sharp dips and steps are seen in several of the subplots, and a very interesting

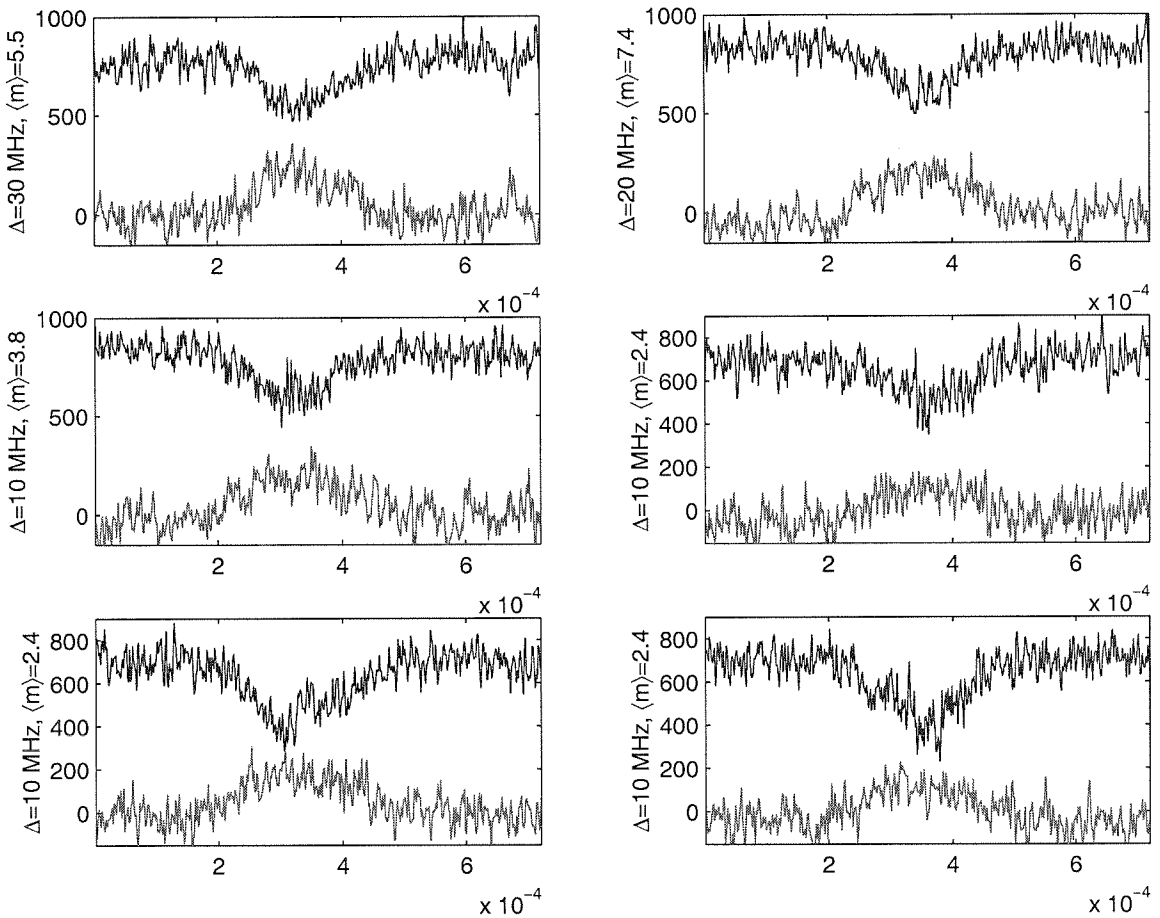

Figure 2.18: Some single-atom transit signals selected specifically because they seem to show some internal structure.

transition from a smooth signal to oscillations can be seen in the middle left subplot. Of course, some of these features are not that much larger than noise events seen in the baseline to either side of the transit signals.

The one thing I feel we can definitely conclude from these data is that our experimental transit signals do achieve the signal-to-noise ratios that they should. Note that the size of the expected transit-induced changes in the phase and quadrature amplitudes is a strong function of the intracavity photon number $\langle m \rangle$. Our experimental calibration of $\langle m \rangle$ was based on a comparison of the mean value of the amplitude quadrature with no atoms in the cavity to the heterodyne shot noise, via equation (2.15). This value was then used to fix $\mathcal{E}$ in the Master Equation (2.3) to produce the simulated signals shown in Figures 2.10 and 2.11. So the fact that the magnitudes

of the observed signals match our theoretical expectations (further evidence for this will be presented below) closes the loop on our overall characterization of η, $\langle m \rangle$, and the signal-to-noise ratio.

I would also say that the experimental signals really don't seem to be undersampled, in spite of what the numerical simulations seemed to suggest. Taken in conjunction with our independent assessments of the technical noise (discussed above), I feel these data demonstrate that we have in fact reached the shot-noise limit for observing time-dependence of the atom-cavity coupling in individual transit events.

2.5.2 Transit phasors

I would also claim that certain aspects of the data show surprisingly good agreement with theoretical predictions based on the cavity-QED Master Equation (2.3). The agreement may be considered surprising because the equation I wrote ignores birefringence of the physics cavity. This means in principle that intracavity atoms interact strongly with *two* electromagnetic field modes rather than one. Also, it seems that optical pumping and Zeeman coherences somehow conspire to make it possible for the atom-cavity coupling strength g to reach its maximal value of 11 MHz, which one would expect should only apply if we were driving the circularly-polarized $6S_{1/2}\,(F = 4, m_F = 4) \rightarrow 6P_{3/2}\,(F' = 5, m_{F'} = 5)$ Zeeman transition. Circular polarizations are not generally supported by our cavity, and the maximal g that would naively be possible for *pi*-transitions should be 6 MHz. By the time one includes the effect of residual magnetic fields and two linearly-polarized cavity modes, however, as in equation (2.12), the situation becomes very difficult to think about in familiar terms. In principle it would be nice to try to simulate transit signals using the full theory, but for the long term I think our efforts would better be spent getting rid of the cavity birefringence altogether.

In Figures 2.19 through 2.21 I display single-transit data on the complex plane. The basic idea is to "put the pen down" just as an atom enters the cavity, and then continuously to trace out the curve of signal amplitude versus signal phase until the

atom leaves. One can clearly identify the appropriate endpoints by looking at the data as in Figure 2.18, and these transit phasors simply represent a way of looking at the correlation between amplitude and phase with time removed from the picture. I believe that this type of plot is the best format for comparisons to theory, because all of the atomic internal (optical pumping among Zeeman states) and external (motion through the cavity eigenmode) dynamics should be factored out. In the simplest approximation, these two types of processes merely induce fluctuations in the atom-cavity coupling g, which would only cause our hypothetical pen to move in and out along a single parametric curve in the complex plane. The overall shape of this curve should be dictated by the interaction Hamiltonian for the atom and cavity mode(s), and is therefore quite easy to compute.

Figure 2.19 specifically shows transit phasors taken for positive, zero, and negative detunings of the probe and cavity, relative to atomic resonance. Recall that Δ is defined as the atomic frequency minus the probe frequency, so negative Δ means that the probe is to the *blue* of atomic resonance. One clearly sees that the phase of the atom-cavity response to an optical driving field changes sign with the detuning, as it should. On resonance, there is no swing in the phase of the transmitted field.

In Figure 2.20, I show a set of transit phasors all taken with $\Delta = 10$ MHz, but with varying strengths of the driving field. Starting from the left, the first column shows data, the second column shows theoretical phasors based on the Master Equation (2.3), the third column shows theory based on the semiclassical equation (2.6), and the fourth column again shows data. I think it's fair to say that the data show some preference for the quantum theory over the semiclassical theory. Note that the "width" of the points in the theory plots is not actually meant to reflect computed variances, but rather just to make the lines more visible. The two experimental signals shown per value of the driving field were selected from the data sets because they were the largest signals at the given detuning and probe strength. So although we are certainly in the low-statistics regime, I would claim that the signal selection was relatively unbiased.

In Figure 2.21 I show a set of transit phasors taken at different probe detunings.

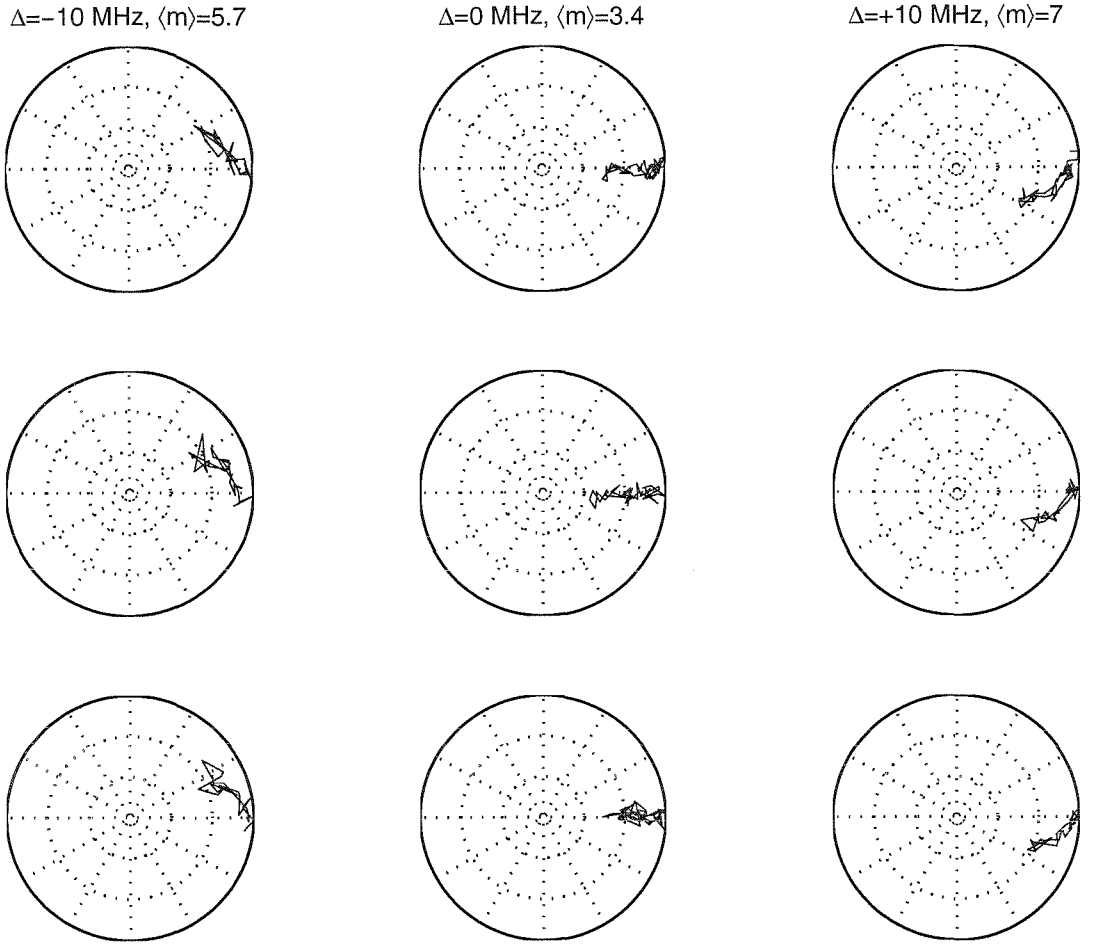

Figure 2.19: Transit phasors showing the general trend from positive ($\Delta \equiv \nu_a - \nu_p = -10$ MHz), to zero, to negative ($\Delta = +10$) detuning of the probe frequency relative to atomic resonance. The data were sampled at 10 MHz, with 12 bit resolution, and then digitally filtered down to a bandwidth of 50 kHz for display.

The leftmost column shows quantum theory, and the other two show data selected in the same fashion as for Figure 2.20. Again, I think the agreement is quite encouraging.

2.6 Discussion: significance and future goals

My primary conclusion about this experimental work is that we have demonstrated the ability to monitor dynamical variations in the atom-cavity coupling strength during individual transit events, with nearly shot-noise limited sensitivity. This claim is supported primarily by a direct comparison of the signal-to-noise ratio observed

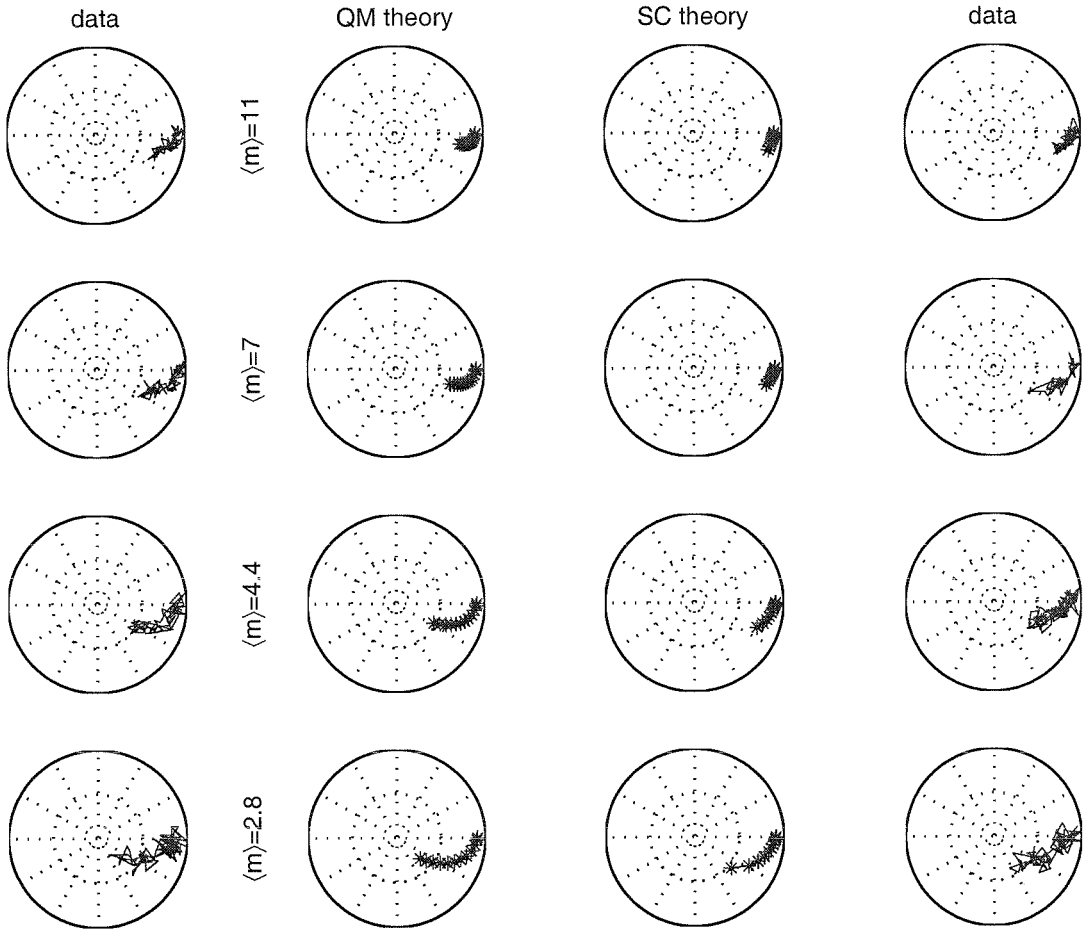

Figure 2.20: Transit phasors for fixed detuning $\Delta = 10$ MHz, with variable probe strength (as indicated). The data were digitially filtered down to a bandwidth of 100 kHz.

in our data with that produced in numerical simulations of the experiment. Looking at the comparison of data against theory with respect to the transit phasors, it also certainly seems that the signals are as large as allowed by the (idealized) Master Equation model of equation (2.3). Our independent characterization of the noise levels also suggests that we are very close to being shot-noise limited in our detection of the transmitted probe beam as well, at least in the frequency range of interest above several kHz. So the signal is what it should be, the noise is what it should be, and I conclude that we have finally succeeded in our goal of suppressing the manifold sources of technical noise in our experiment down to the level of the fundamental

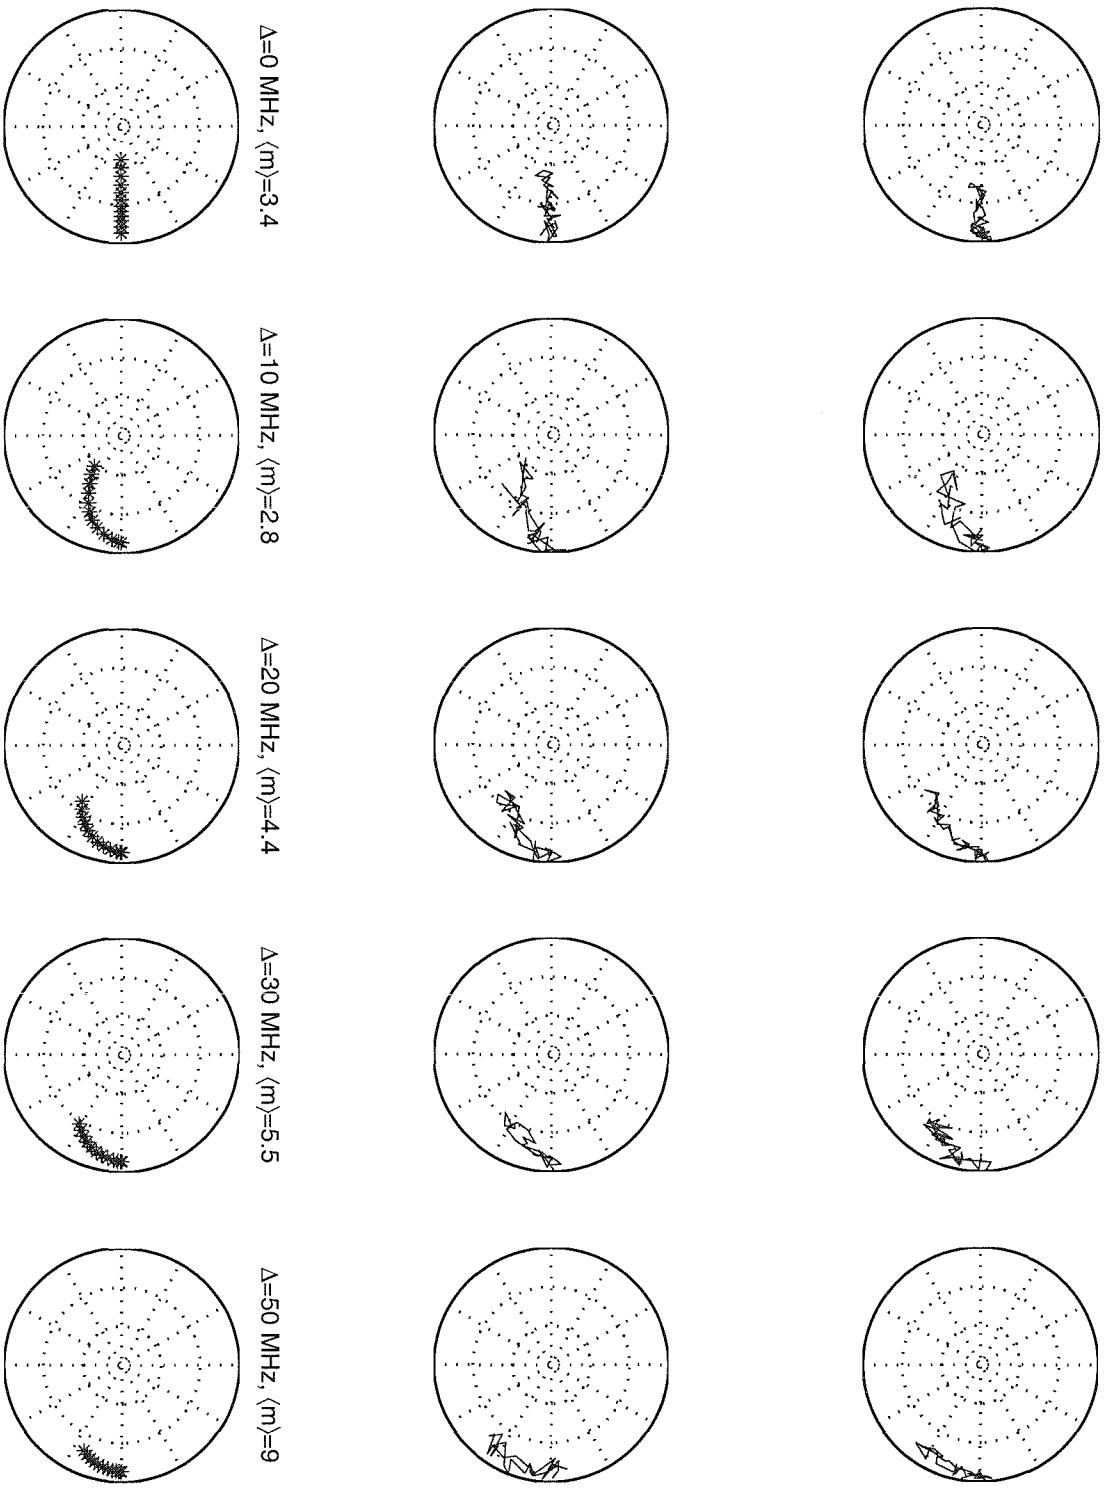

Figure 2.21: Dependence of transit phasor shapes on detuning (as indicated). The left column shows quantum theory, the two other columns show data filtered to a bandwidth of 100 kHz.

vacuum-fluctuation limit for balanced heterodyne detection.

One should appreciate, of course, that all these statements are made relative to the overall photodetection efficiency that we achieved in the experiment. I think that our figure of $\eta = 0.32$ qualifies as being basically respectable, although it should really be possible to reach 0.9 or greater with better photodetectors and a much more serious effort at optimizing the spatial overlap between the signal beam and the optical local oscillator. In order to be really serious about pursuing quantitative tests of conditional quantum dynamics, we would also need to use a single-sided (rather than symmetric) cavity.

Certainly the most significant result that *could* have come of this work, however, would be some unambiguous evidence that we could track the motion of an individual atom along the intracavity standing-wave. Based on a comparison of the numerical simulations and our data, I do believe that some of the more conspicuous transient features we observe in our signals may be attributed to atomic motion along the standing wave. Again, however, there's no way to be sure in the current incarnation of the experiment.

We can look at signals such as those in Figure 2.17 and estimate the sensitivity of our measurements to variations in g. These plots are made at a bandwidth of 300 kHz, and it would seem that the full-signal to rms-noise ratio in each quadrature amplitude is about 5. This would set the sensitivity to fractional variations in g at about $(3 \times 10^{-4})/\sqrt{\text{Hz}}$. In principle, the data should therefore show an optimal sensitivity $\sim 10^{-10}$ m$/\sqrt{\text{Hz}}$ to atomic displacements along the standing wave. As will be shown in Chapter 4, the "standard quantum limit" for overall observation time of a single atom's position with this sensitivity should be on the order of 10 μs for a signal bandwidth of 300 kHz. Recall that the characteristic width (in time) of our transit signals is about 250 μs. What this means is that even with the most highly-idealized, canonical, broadband position measurement allowed by quantum mechanics, we would *necessarily* have to see *some heating* of the atomic motion (due to measurement backaction) if we maintained a sensitivity of 10^{-10} m$/\sqrt{\text{Hz}}$ for a total observation time longer than 10 μs.

But what does this imply about our actual experiment? Independent of any considerations of fundamental measurement backaction, it seems that we do induce some excess heating of the atomic motion due to dipole-force fluctuations and spontaneous-emission recoils. We also can't say that we're really observing the atomic motion in our data, due to obfuscation by atomic internal-state dynamics and motion transverse to the standing-wave axis.

My opinion is the following. By achieving as high a sensitivity to atomic motion as we might claim over the total observation times that we certainly do have, we have pushed our experiments on cavity QED into a regime where it *should* be possible to study conditional quantum dynamics of the sort discussed in the introduction. But because we had insufficient control over the initial positions and velocities of the incoming atoms, and because we could not restrain the atomic internal dynamics to a simple two-level manifold, the equation we would have to use in conducting such studies is just ridiculously complicated (*c.f.* equation (2.12)).

In order to get back into a reasonable regime for scientific inquiry, I think it's clear that we need to do two things. First, we have to learn to trap and cool individual atoms inside the cavity to prepare well-defined initial states of motion. Second, we need to make cavities with no significant birefringence, and drive the atom-cavity system with light of very pure circular polarization. Even better, we should try to find an atom other than Cesium that has a $J = 0$ ground state! It might also be wise to try a lighter atom, in the sense that life would be a lot easier if we could hold our measurement sensitivities fixed but increase the magnitude of the backaction we wish to study.

In any case, however, I do feel that we have achieved a significant technical advance in our ability to perform *precision* measurement in the context of cavity quantum electrodynamics. I have no doubt that this will enable lots of great physics that seemed to be just out of reach several years ago—not only in the field of conditional quantum dynamics, but for realizing experimental proposals for quantum logic, quantum communication, and perhaps the observation of single-atom optical bistability as well.

Chapter 3 Technical commentary

My primary purpose in this chapter is to fill in further technical details of exactly how we did things. I have also tried to weave in a few historical comments on the many rebuildings and overhauls that "the" experiment has undergone in the years since my very first efforts with Quentin. As I would assume must be the case with most experiments that evolve in this fashion, many aspects of what we have today were determined more by legacy than by careful engineering. Still, I hope that I have managed to convey some sense of why the apparatus has grown to have the form that it does, and where possible to provide some record of the things that were tried early on but eventually proved to be inadequate. I have also taken this Chapter as an opportunity to try to convey some of the lessons I think I've learned during more than a few hours invested in serial trial-and-error. I hope that someone may find something of value as a result, and that everyone else will forgive the didactic urges that occasionally possess the objective description of details.

3.1 The grand tour

In terms of hardware, the apparatus has a number of interconnected parts (see Figure 3.1). There are four lasers, three high-finesse cavities, seven major servo-loops, and a vacuum system. At the center of it all is what I'll call the *physics cavity*, so let's take this as our point of departure for a walk-through tour of the experiment.

3.1.1 The physics cavity

The physics cavity is where we would like cavity-quantum electrodynamic interactions to take place between individual atoms and just a few photons. Hence it has small mode-volume, high finesse, and lives under vacuum. One particular cavity with 106

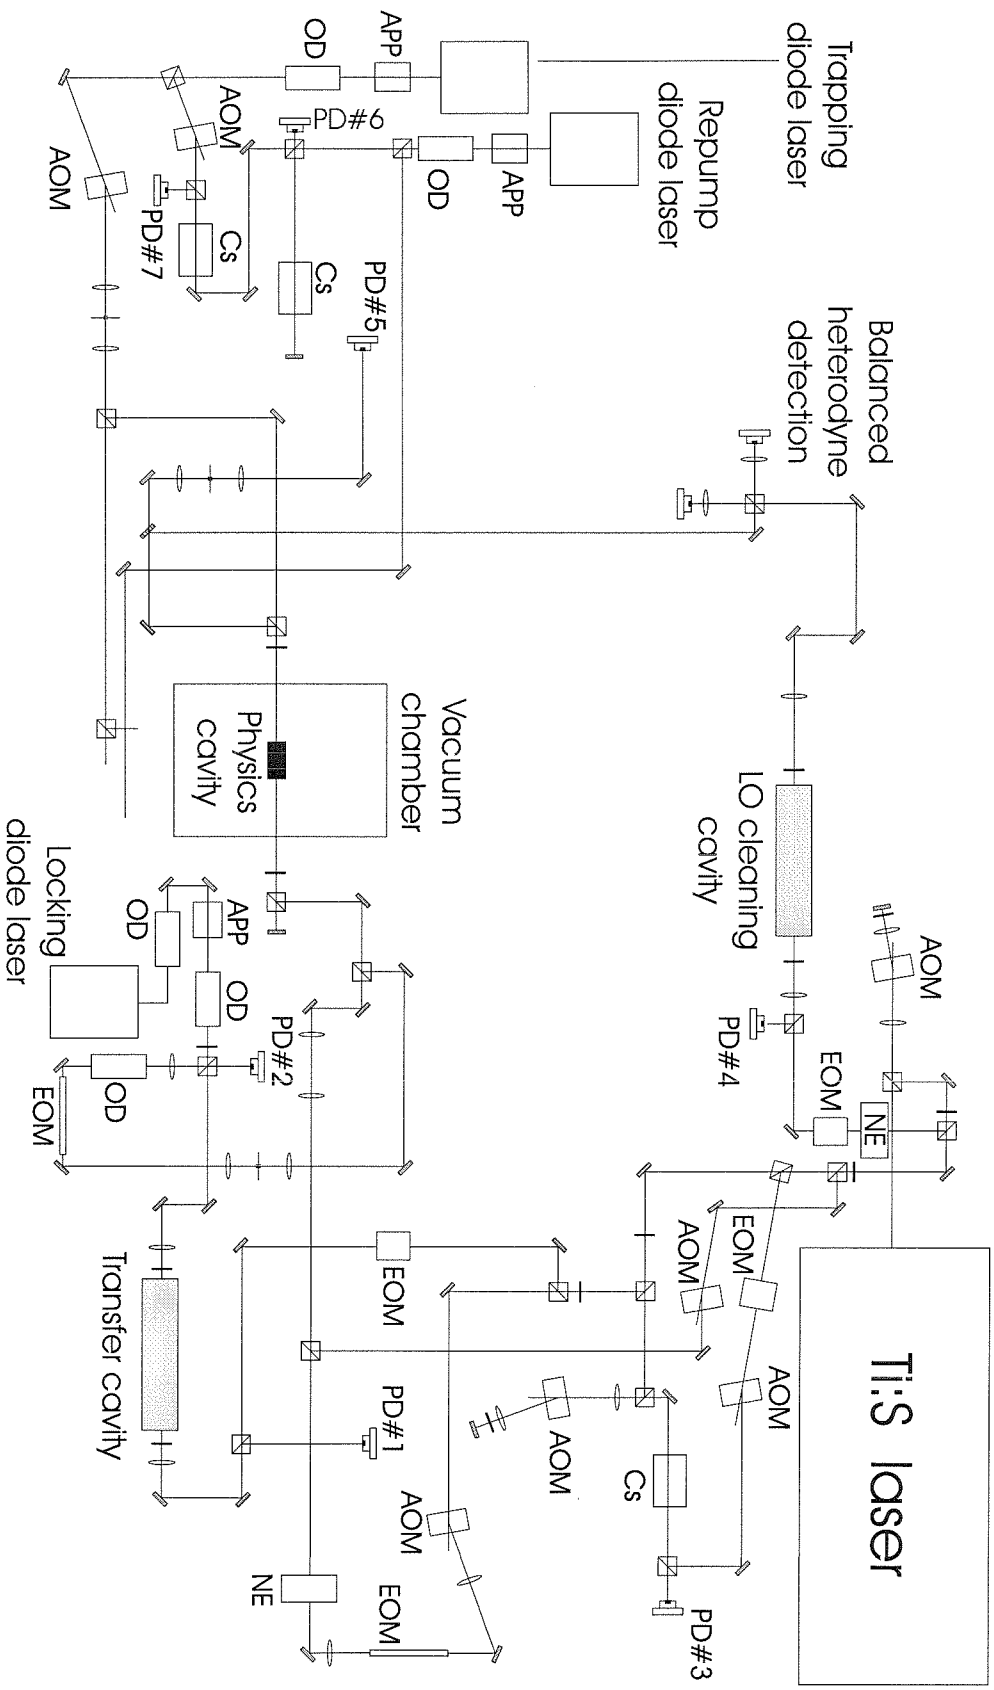

AOM - acousto-optic modulator
EOM - electro-optic modulator
NE - noise eater (intensity stabilizer)

Cs - cesium vapor cell
APP - anamorphic prism pair
OD - optical diode (isolator)

PD#1 - Ti:S lock to transfer cav
PD#2 - Diode lock to transfer cav
PD#3 - Transfer cav lock to Cs
PD#4 - Cleaning cav lock to Ti:S

PD#5 - Physics cav lock to diode
PD#6 - Repump lock to Cs
PD#7 - Trapping lock to Cs

Figure 3.1: Detailed layout of the optical table. Some steering optics and diagnostic paths not shown.

μm length and $\mathcal{F} \simeq 212,000$ was used for all of the measurements described in this thesis, although in retrospect we probably should have changed to a new one long ago. This cavity has certainly treated us well, but it suffers from a pronounced birefringence that completely foiled our attempts at cleanly isolating strong σ_+ transitions within the Cesium D2 Zeeman manifold.

The choice of length $\sim$ 100 μm was intentionally rather conservative (following the 50 μm "one-dimensional atoms" cavity [18]), to the extent of taking a rare step backwards in the Kimble Group's march towards ever smaller critical atom/photon numbers. In fact this was the third cavity that we tried, having first gone through an old-style cavity (in one of the big brass and invar mounts [19]) and a cavity with 3 mm-diameter substrates glued to a very fancy "monolithic" mount inspired by those EDM'ed Physik Instrumente piezostages. We never saw any signals with either of the first two cavities, and the design of the third was dictated by a play-it-safe strategy for eliminating all the cavity-related obstacles we could think of.

Our initial experience with the old-style cavity taught us that we had to eliminate as much metal as possible from the mount itself—scattered light from the MOT beams was a serious nuiscence, and we could never get a trap to form anywhere close to the cavity axis. It had also been clear to us that the standard 7 mm substrates were just too big, but I think that at the time we were still waiting for Research Electrooptics to turn down smaller ones. When the small substrates finally came, Quentin designed an "open" monolithic mount that allowed us to bring the trapping beams right down to the mirror substrates. This incorporated a plate PZT (as opposed to the traditional tubes) along with some EDM'ed leaf springs to actuate the cavity length from below. We never quite pinpointed the problem with this cavity, and in retrospect, it may not have actually had any serious ones. We identified two potential problems at the time: first, it was made from somewhat-magnetic 304 stainless steel, causing us to fear that the falling atoms were somehow being deflected in residual magnetic fields even after the MOT coils were turned off; second, the cavity was quite short (10 μm, I believe) but had 3 mm-diameter substrates. Judging from what we and Hood/Chapman/Lynn have done since, I'd guess that neither of these flaws were truly fatal. However, that

was certainly not the right cavity to start with.

Just to save the information, let me note that the electron-discharge machining was done quite nicely by Bangs Manufacturing of Burbank, CA (818-845-3528). They had previously made some translation stages for Olivier Carnal, which were used to position the gratings in the OLGA beam machine. Overall, the plate PZT/EDM'ed springs design seems to be fine for DC positioning, but at least our mount had rather low resonances and was mechanically pretty noisy. One should rather go the other way and try to mount a low-mass mirror substrate directly onto a tweeter PZT to get the fastest possible response, which Dave and Jun will presumably try to do for the next cavity design.

As I recall, we had hoped to end up with a somewhat shorter cavity than 100 μm, but when the glue dried that was what it was. The first things we did once the present cavity was aligned in the vacuum chamber and pumped down were to measure its length and finesse. At that point it immediately became clear that we had a problem with birefringence, but we pressed on anyway. One significant point to note here is that we have not seen any degradation whatsoever in the cavity finesse during more than two years of service. This may possibly be attributed to the fact that we have always run a beam-loaded MOT, rather than a vapor-loaded MOT as in Hood/Chapman/Lynn or a direct thermal beam as in previous experiments. Our experience of not being able to control the cavity length and the evident problem with birefringence have motivated Christina and Mike to develop their high-tech cavity assembly method and jig. Given the success of their 10 μm cavity, we may hope that the problems Quentin and I initially encountered have truly become a thing of the past.

3.1.2 Locking the physics cavity

It should come as no surprise that locking the physics cavity *well* has been a major technical challenge. Quentin, Mike and I started looking for atom transits with an old-style cavity lock, by dithering the physics cavity PZT somewhere between 30-80

kHz and using a strong lock beam at 50% duty-cycle [18]. As opposed to the fast electrooptic chop that Christina and company have been using of late, back then we used a mechanical chopper in the lock-beam path that only got us up to 4 kHz or so. A piece-of-glass type beamsplitter sent some transmitted light to a PMT, whose photocurrent was directed to a lock-in amplifier to generate an error signal for the physics cavity PZT. That lock was terrible.

Substantial improvement occurred when Quentin realized that we could get 100% locking duty cycle by sifting an error signal out of the *heterodyne* photocurrent. This meant that we no longer needed a super-strong lock beam, at the cost of having to turn up the cavity dither and of sacrificing our ability to run the experiment at very low probe powers. The heterodyne/cavity-dither method got us through the era of the 1996 Optics Letter, but was always a severely limiting factor for the overall quality of the data. Its major drawbacks were inherently low signal to noise (probe powers out the back of the cavity are typically $\sim$ picowatts), extreme nonlinearity (as input to the lock-in we just took the logarithmic "Video Out" signal from the spectrum analyzer), and strong cavity dither (which certainly showed up to some degree in the data). Christina, Mike, and Theresa have implemented a significantly-improved version of this locking method, but as my interests drove our experiment towards the precision-measurement regime it quickly became clear that something drastically different was called for.

In order to get an error signal with 100% duty cycle and good signal-to-noise at high bandwidth, it seems that one has to use light that is far-detuned from Cesium resonances. The natural scheme is to hop one or more free spectral ranges away from 852.359 nm and lock the physics cavity there. In terms of Cesium transitions, the D2 is bracketed on either side by the D1 at 894 nm and the $6P_{3/2} \rightarrow 8S_{1/2}$ at 794 nm. You want to stay away from all three of these lines, which puts you either around 825 nm or 873 nm. For reasons of detector efficiency, one would rather go towards shorter wavelengths, but I found that the 15 ppm coatings of our physics cavity rolled off fairly steeply starting around 820 nm.

Two major technical considerations complicate this scheme: first, you clearly

need another (diode) laser to provide light that is sufficiently far detuned from the Ti:Sapphire at 852 nm, and second, the modes of a 100 μm cavity are so sparse that you need to find some way of locking a laser out in the middle of nowhere (with respect to convenient atomic references, that is). A search for available diode lasers revealed that an SDL model at 810 nm was about the best we could do, although I would have rather stayed closer to 852 nm in order to maintain higher cavity finesse for producing an error signal. In terms of finding something to lock to, I could only think of two possibilities.

The first idea I investigated, but eventually abandoned, was to try to find a reasonable set of molecular lines that would be dense enough around 810 nm to use an rf-offset lock for stabilizing the aforementioned SDL diode laser. It turns out that Iodine Monobromide (IBr) has lots of known transitions in that wavelenth range [60, 61], but making an IBr cell was sufficiently nontrivial (IBr easily dissociates into I_2 and Br_2 upon heating) and the molecular lines sufficiently weak that I decided to try a different approach first.

That was to use an auxilliary ("transfer") cavity for relative stabilization of the 810 nm diode to the Ti:Sapphire. I found some miscellaneous mirrors from Eugene Polzik's stash, which had about the right transmission and curvature and were sufficiently broad to support high-finesse modes at both 850 and 810 nm. Quentin and I assembled this into a fairly long cavity using one of the Texas-era brass/invar housings, with the $\sim$ 15 cm length determined by the need to have a mode spacing of $\sim$ 300 MHz or less. This constraint comes from the fact that the free spectral ranges of the physics cavity and transfer cavity are incommensurate, so you have to be able to find a transfer-cavity mode within electrooptic reach of some physics-cavity mode in the neighborhood of 810 nm (the physics cavity FSR is about 3.3 nm). We use one of the Mark Raizen travelling-wave modulators to do this. One point to note here, however, is that we found we had to use the TWM in the path to the physics cavity rather than the transfer cavity—doing things the other way around you just get too many resonances since the TWM sideband spacing is comparable to the transfer cavity mode spacing.

3.1.2.1 Locking the cavity-lock diode laser

I put together a Libbrecht-design current controller and laser housing for the SDL 810 nm diode, with a current-dither around 8 MHz to generate both the diode-to-transfer-cavity and physics-cavity-to-diode FM error signals. In early days the diode locked to the transfer cavity extremely well, but I suspect that this should be attributed to some fortuitous optical feedback of just the right strength to help out my electrical servos.

At the time we first tried the transfer-cavity scheme there was no room anywhere on the optical table to put all this stuff, so it all went on a breadboard that we mounted above the trapping diodes. This placement caused great inconvenience when the diode wavelength drifted and I had to tweak the grating. Over a 4 m pathlength (with a pinhole spatial filter and travelling wave modulator in the way), such adjustments caused serious misalignments of the beam going to the physics cavity. So after Quentin left to write his thesis, I decided to compactify the setup and moved everything down as close to the physics cavity as I could manage.

This seemingly innocent move aggravated a number of problems with the SDL diode laser setup. The worst problem was the fact that the diode laser now sat directly under the main air-conditioning vent, which blasts very cold air down onto the optical table. This caused violent excursions (100's of MHz peak-to-peak) of the diode frequency, with correlation times on the order of 10 ms. There's basically nothing you can do to lock those out with an electronic servo. Although one might think that a hermetic sealing of the diode laser housing should fix such problems, I have found that the frequency of SDL diode lasers with Littrow grating feedback is quite sensitive to any sort of pathlength jitter along the first 50-100 cm beyond the laser housing. This is presumably due to residual optical feedback, and in any case it makes such diodes extremely intolerant of turbulent air on the table. Rather than attempt to enclose all the diode optics, I opted to redirect the flow of air in the lab. Curtains around the table could probably also have done the trick, but they're hard to work with on a day-to-day basis.

Somehow during the relocation of diode laser and transfer cavity, locking conditions for the diode laser became much less favorable in general, and the diode-to-transfer-cavity lock was no longer good to better than the transfer cavity FWHM. I was using a two-branch servo scheme, with feedback to both the grating PZT (unity-gain at 10 Hz) and the diode current (unity-gain at 6 kHz). The overall behavior of the lock was rather strange—the laser would basically lock cleanly to a transfer cavity resonance (to better than 5% of the cavity's 1 MHz FWHM), but would burst into irregular high-frequency (10's of kHz) oscillations for several ms at a time with 10-100 ms in between. These intermittent bursts would have peak amplitudes of something like the transfer cavity FWHM, and the physics cavity seemed to follow their $\sim$ 10 Hz envelope. So this was a real problem. I eventually decided to buy a New Focus diode laser system (which we run at 836 nm) to replace the ill-behaved home job, and although that laser performs beautifully in general, it exhibits exactly the same sort of behavior when its grating starts to creep out of alignment. Hence, my understanding is that the poor performance of the SDL diode was due to competition between the optical feedback from its grating and from the transfer cavity. In the New Focus laser, which is AR coated, the grating feedback is strong enough (when optimized) to win that battle outright and keep the diode frequency pinned down. Of course, an alternative way to go would be to set up a pathlength servo to try and let the optical feedback from a cavity control the diode frequency instead—the New Focus lasers are rather expensive for what they are, and it's not clear they're going to continue carrying them as a commercial product.

3.1.2.2 The physics cavity itself

The physics cavity mount that we have been using since our very first atom transit signals was designed and assembled by Mike Chapman. A drawing was included in Quentin's thesis. The actuator for the cavity length is a 1 cm diameter, 50 mil wall thickness tube PZT, with some amount of precompression by a stiff metal spring. When I set out to design a more serious servo for the physics cavity, my first step was to try to measure the open-loop transfer function from PZT voltage to mirror

displacement.

Unfortunately the open mount structure and small-footprint vibration isolation stack that we have had to use are insufficient to maintain passive stability of the cavity resonance within its narrow ($\mathcal{F} \sim 212,000$) linewidth. So I had to establish a weak lock of the physics cavity to the 810 nm diode laser first. The idea here was to try to close the loop at some very low frequency (like 50 Hz), so that the response up in the kilohertz regime would be more-or-less unaffected by the servo. Assuming that the open-loop transfer function should be pretty flat at frequencies below 1 kHz, you'd get all the necessary information to design a more serious controller that closed as high as the open-loop response would allow. The test controller I put together was a simple integrator (300 kΩ and 1 μF around an OP27), plus an extra op-amp to sum in sinusoidal drives from the HP Dynamic Signal Analyzer (DSA).

As it turned out, there was no way to get the cavity even vaguely locked without turning up the gain until the unity-crossing point was way up at 885 Hz. Undaunted by this, I recorded the closed-loop response with some modest degree of averaging (shown in Figure 3.2). I eventually identified the mini-features below 10 kHz with noise peaks of the 810 nm diode laser, so it appeared that the first real catastrophes in the physics cavity response occurred at 10 kHz (the pi phase swing) and 14 kHz (the largish amplitude peak). Something funny was also going on with the phase around the unity-crossing point of the servo, but I decided to write this off as a purely closed-loop effect to begin with. Given the major features above 10 kHz, I thought 1 kHz would be a reasonable target for closing an aggressively-designed physics cavity servo.

I built a PI controller with low-pass corners in the proportional branch at 3 Hz and 10 Hz. It hit a minimum phase of -160 degrees at 70 Hz, and had 100 dB attenuation and -135 degrees of phase at 1 kHz. It didn't work.

The problem seemed to be unity-gain oscillations, even though I had designed a comfortable phase margin of 45 degrees. Looking back at the response measurement, the closed-loop transfer function was certainly doing nontrivial things by around 3 kHz, so it seemed like it might be worth the trouble to try and divide out the controller action to infer the physics cavity's open-loop response more precisely over as

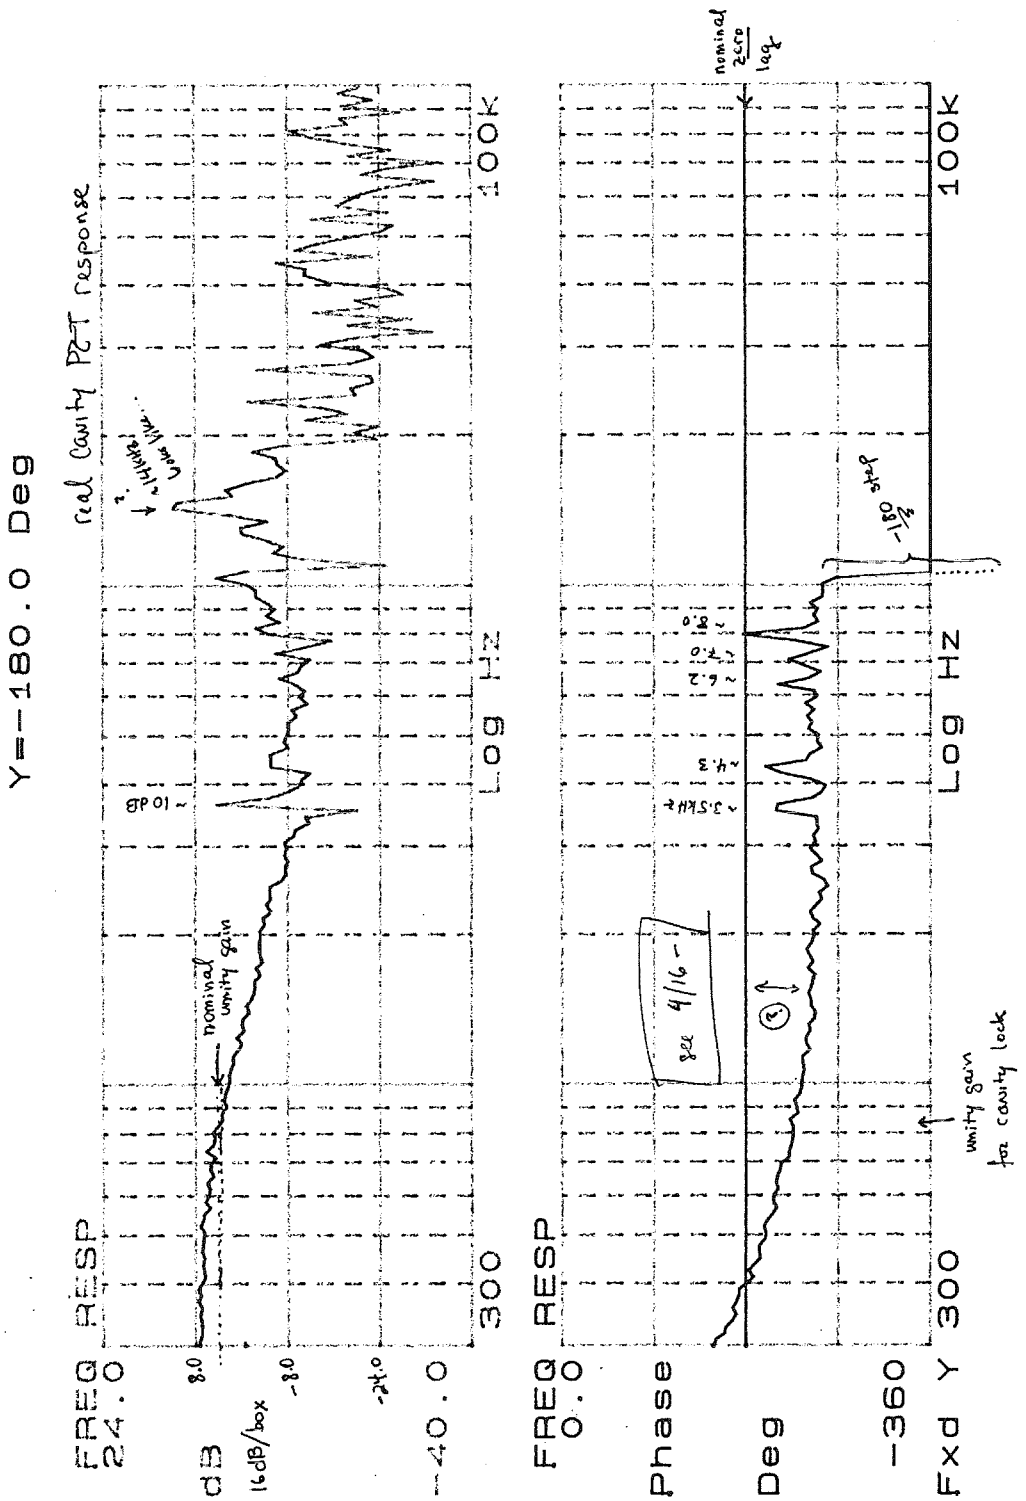

Figure 3.2: Closed-loop response of the physics cavity with test servo.

wide a frequency range as possible. With some patience you can use the HP DSA (in swept-sine and auto-integrate mode) to accumulate *very* clean transfer-function measurements, and GPIB allows you to transfer these to a PC for convenient processing with Matlab.

The standard Laplace-domain treatment of linear feedback control tells you that the closed loop transfer function $\tilde{G}(s)$, open-loop transfer function $G(s)$, and controller transfer function $H(s)$ are related by

$$\tilde{G}(s) = \frac{G(s)}{1 + G(s)H(s)}. \tag{3.1}$$

So having measured any two of these you can infer the third. Figure 3.3 shows the measured $\tilde{G}(s)$ and $H(s)$ for the physics cavity and test controller, together with the inferred $G(s)$. My understanding is that the feature around 50 Hz is not actually a mount resonance, but rather a noise peak due to some resonant mode of the vibration isolation stack. The pertinent feature of $G(s)$ is really the $\approx$ 30-degree phase lag around 1 kHz, which was clearly large enough to destabilize the fancy controller. If you have the signal-to-noise in your error signal such a problem should be quite easy to fix, so I added a phase lead network with a zero at 500 Hz and a pole at 5 kHz on the output side of the fancy controller. Much to my amazement, this attempt at an intelligent fix worked quite nicely. Figure 3.4 shows the original controller transfer function, the effect on it of the excess cavity phase lag, and the corrected controller TF.

Despite this small triumph of local servo wisdom, we eventually replaced the fancy controller with a JILA loop filter that Jun put together, since that design has knobs on the outside of the box for adjusting all the gains and corners. It also features a much cleaner, low-noise construction, which turns out to be quite important for minimizing excess noise on the cavity lock. As far as we have been able to tell, the dominant remaining problem with the cavity lock these days is intermittent oscillation of those 10 and 14 kHz resonances. These oscillations seem to be electrically driven, as we have found that we can kill one or the other of them by putting LC filters on

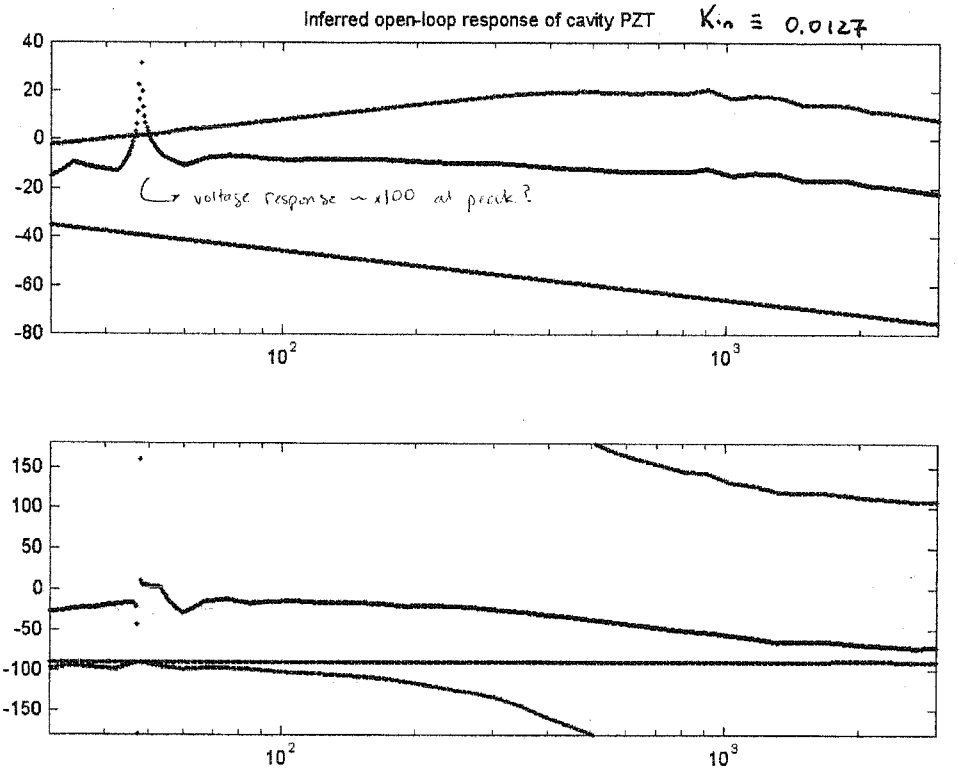

Figure 3.3: Closed-loop, controller, and inferred open-loop transfer functions. In the upper subplot, which shows magnitude, the lowest trace shows controller response, the top trace shows the closed-loop response, and the middle trace with the peak shows the inferred open-loop response. The lower subplot shows phase, where the trace with what looks like a dispersive resonace is the inferred open-loop response, the trace which wraps around from -180 to +180 degrees at $\sim$ 500 Hz is the measured closed-loop response, and the flat line at $-90°$ is the controller response.

the voltage going to the physics cavity PZT. Unfortunately we can't seem to get them both without also eating up all our phase margin at 1 kHz.

I'd finally like to note that the switching of the magnetic field coils of the MOT somehow kicks the cavity lock in a fairly serious way. It's not completely clear whether this is some kind of inductive pickup, or an actual mechanical disturbance. In any case, I would strongly advocate minimizing the magnetic coil size in future designs, as well as mounting them in some manner that keeps them mechanically isolated from the vacuum chamber.

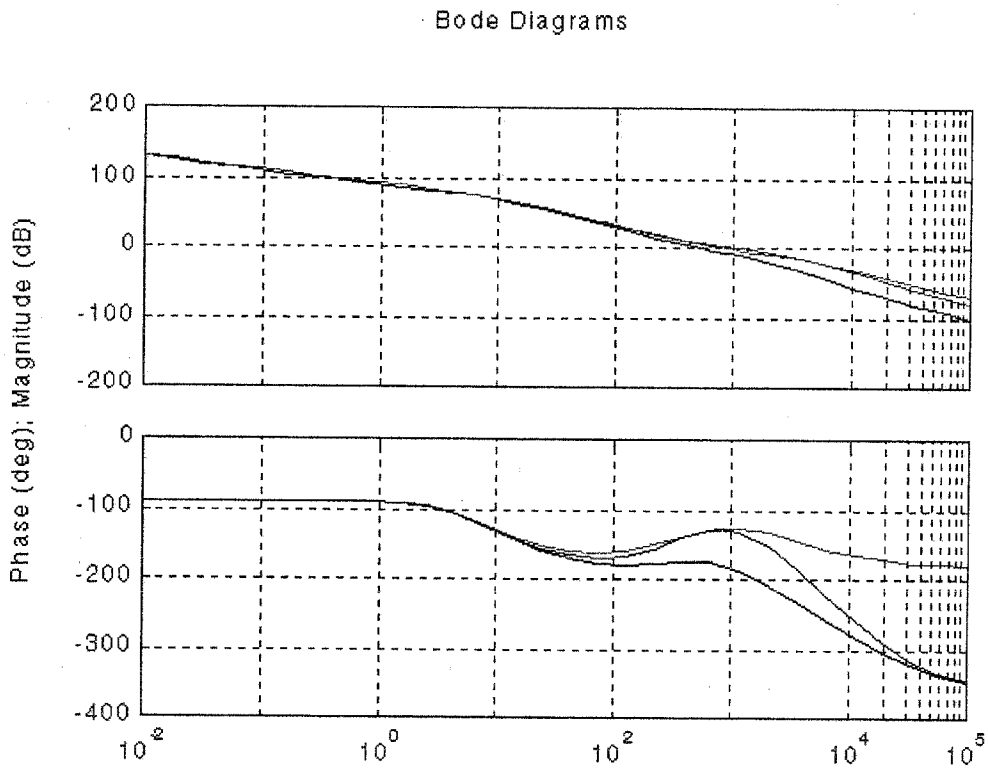

Figure 3.4: Original, perturbed, and corrected servo transfer functions. Again, the upper subplot shows magnitude and the lower subplot shows phase. The top trace in each subplot shows the intended servo design, the bottom trace shows the "actual" response due to PZT rolloff, and the middle trace shows the corrected transfer function after adding an additional phase-lead network.

3.1.3 Locking the transfer cavity

Our original grand strategy for frequency stabilization involved locking the Ti:Sapphire directly to Cesium saturated absorption signals, and then locking the transfer cavity to the Ti:Sapph. This was doubly wrong-headed, as the Ti:Sapph stability is much improved by locking to the high-contrast cavity resonances and it was basically impossible to lock the transfer cavity to the Ti:Sapph anyway. As I'll describe in more detail below, our Coherent 899 has terrible passive frequency stability, exhibiting several-MHz excursions on 1-100 ms timescales. The transfer cavity linewidth at 852 nm is around 100 kHz. Even when it was nominally locked to Cesium, the Ti:Sapph frequency would whip through the cavity FWHM in ~ 0.1 ms as it wandered around in frequency space, so you just can't get the speed to establish a lock if you're driving

one of those big tube PZT's to actuate the cavity length. In our initial efforts with the transfer cavity, Quentin and I figured that we could count on the cavity to have decent high-frequency stability, and that the main point of locking it to the Ti:Sapph was to prevent slow drift of the cavity length. Hence, we used a lock that we thought should have negligible gain except at DC and ignored the horrible-looking error signal that resulted. In hindsight, I don't think we really knew what the unity-gain point of that servo was, so I wouldn't be too surprised if we were actually transfering Ti:Sapph noise onto the transfer-cavity length at frequencies $\sim 1 - 100$ Hz.

Since Jun brought the acousto-optic laser lock back into style here (for the first time since Tim Boyd's days), we've been locking the Ti:Sapph directly to the transfer cavity. The transfer cavity then gets guided by an error signal from a Cs spectrometer that monitors the absolute frequency of the Ti:Sapph. I've been using a modulation-transfer setup in the Cs spectrometer, both because this configuration eliminates Doppler background from the signal and because the actual $F = 4 \rightarrow F' = 5$ transition is the strongest line (as opposed to the case of FM saturated-absorption spectroscopy, in which the $4, 5$ crossover dominates).

When the Ti:Sapph is locked to the transfer cavity, you can tune the bias on the cavity PZT to scan the Ti:Sapph frequency over the Cesium lines or a resonance of the local oscillator cleaning cavity. Judging from the excellent stability of the trace you see when doing so, the passive stability of the transfer cavity seems to be at least as good as 10 kHz rms in a 100 kHz bandwidth.

3.1.4 The Coherent 899 Ti:Sapphire laser

Although there are many nice things about the Coherent 899-21, it's clear that it really wasn't built to be a narrow-linewidth laser. The spec they give you is something like 1 MHz, and you can tell that they were serious about trading off stability in favor of tunability when they engineered the laser controls and intracavity elements.

First off, the 899-21 incorporates a thick-etalon servo that surpresses the occurrence of mode hops as you tune the brewster plate and/or thin etalon. The thick

etalon is mounted in the ICA (intra-cavity assembly) on some sort of PZT arrangement that permits 2 kHz dithering of the effective length of the thick etalon. This produces a 2 kHz response in the laser *power*, which is demodulated by some electronics in the 899 control box to produce an error signal for stabilizing the angle relative to the thick etalon modes of the laser cavity. Of course, the dithering also induces a 2 kHz modulation of the laser *frequency*. The factory setting for the dither amplitude was such that the frequency excursions were roughly a few hundred kHz, which is small enough to satisfy the Coherent linewidth spec yet (presumably) large enough that the desired error signal has sufficient signal-to-noise ratio even if the laser is operated somewhere in the tails of its tuning range. It turns out that there's a convenient potentiometer to control the dither amplitude, which I identified and marked on the bootleg copy of the control-box schematics we got from a friendly Coherent engineer. With the laser at 850 nm, and with the cavity reasonably clean and well-aligned, the thick etalon servo is in fact okay with a much smaller dither amplitude. Having made this adjustment, the residual frequency modulation at 2 kHz is now small enough that it doesn't poke up in a power spectrum of residual frequency noise of the locked Ti:Sapph.

Another annoying feature of our 899-21 (I'm not sure if this is a design flaw, or just a problem with our particular unit) is that something about either the brewster-plate galvo or the galvo current driver induces an excessive amount of low-frequency noise in the laser frequency. Line-synched glitches can actually be seen in the voltage drop across a fake resistive load in the current loop, although it's not completely clear that they're large enough to account for the problem. My general mistrust of the brewster-plate galvo stems mainly from the fact that it's much easier to lock the Ti:Sapph frequency to a narrow line (like that of the transfer cavity) if you unplug the galvo from the Coherent control box. In fact, I found that I could lock the Ti:Sapph to the transfer cavity using only the tweeter PZT (*i.e.*, without using an external AOM) if and only if the galvo was disconnected. The circuit schematic for the galvo driver indicates some sort of "centering" mechanism that may intentionally tweak the galvo angle at regular intervals (the schematic shows a 20 ms timer), but I haven't

had a chance to figure out what this thing really does.

Moving on to the 899 tweeter, our unit has a much larger (in terms of both size and capacitance) PZT than is being used in home-built Ti:Sapph's around the lab. Figure 3.5 shows a measurement of the PZT response, as determined by driving the PZT with a swept-sine while monitoring an FM saturated-absorption Cesium signal. The dominant features in the transfer function are a clump of resonances around 17 kHz, and in my experience these do set the limit of how well this Ti:Sapph can be stabilized using the tweeter alone. Using a servo that crosses unity-gain at 9 dB/octave and gradually gets back to 12, I've found that you can't really close any higher than 6 kHz without seeing intermittent oscillations. Incidentally, the tweeter PZT leverage for tuning the Ti:Sapph is about 5 MHz per volt.

3.1.5 Locking the Ti:Sapphire

The tweeter servo built into the Coherent control box for the 899 is reasonably sophisticated, but conservative in design. When the 899 first arrived, someone here (probably Oliver Carnal) picked out the wires carrying the error signal intended for locking the Ti:Sapph to its Fabry-Perot cavity and added an "external error signal" input on the back of the control box. Driving this input, and looking at the actual high voltage that would go to the tweeter, you find that the servo amplifier is flat from around 1-10 Hz, rolls off towards DC (where the brewster plate galvo is presumably meant to take over), and maintains a beautiful 9 dB/octave slope all the way out to around 30 kHz. Looking back in my lab book, I'm not positive that I didn't use a signal amplifier for that measurement, but the plot I have shows an open-loop gain of -38 dB and 40 degrees of phase margin at 30 kHz. If we had a better tweeter, it looks like an error signal of 15 volts/MHz would let you close at 30 kHz and get about 110 dB of gain at low frequencies. To stick with a unity-gain of 6 kHz, however, 2.5 volts/MHz is about the right size for the error signal. In order to maintain adequate phase margin, the error signal should not be filtered with a low-pass corner lower than 30 kHz or so.

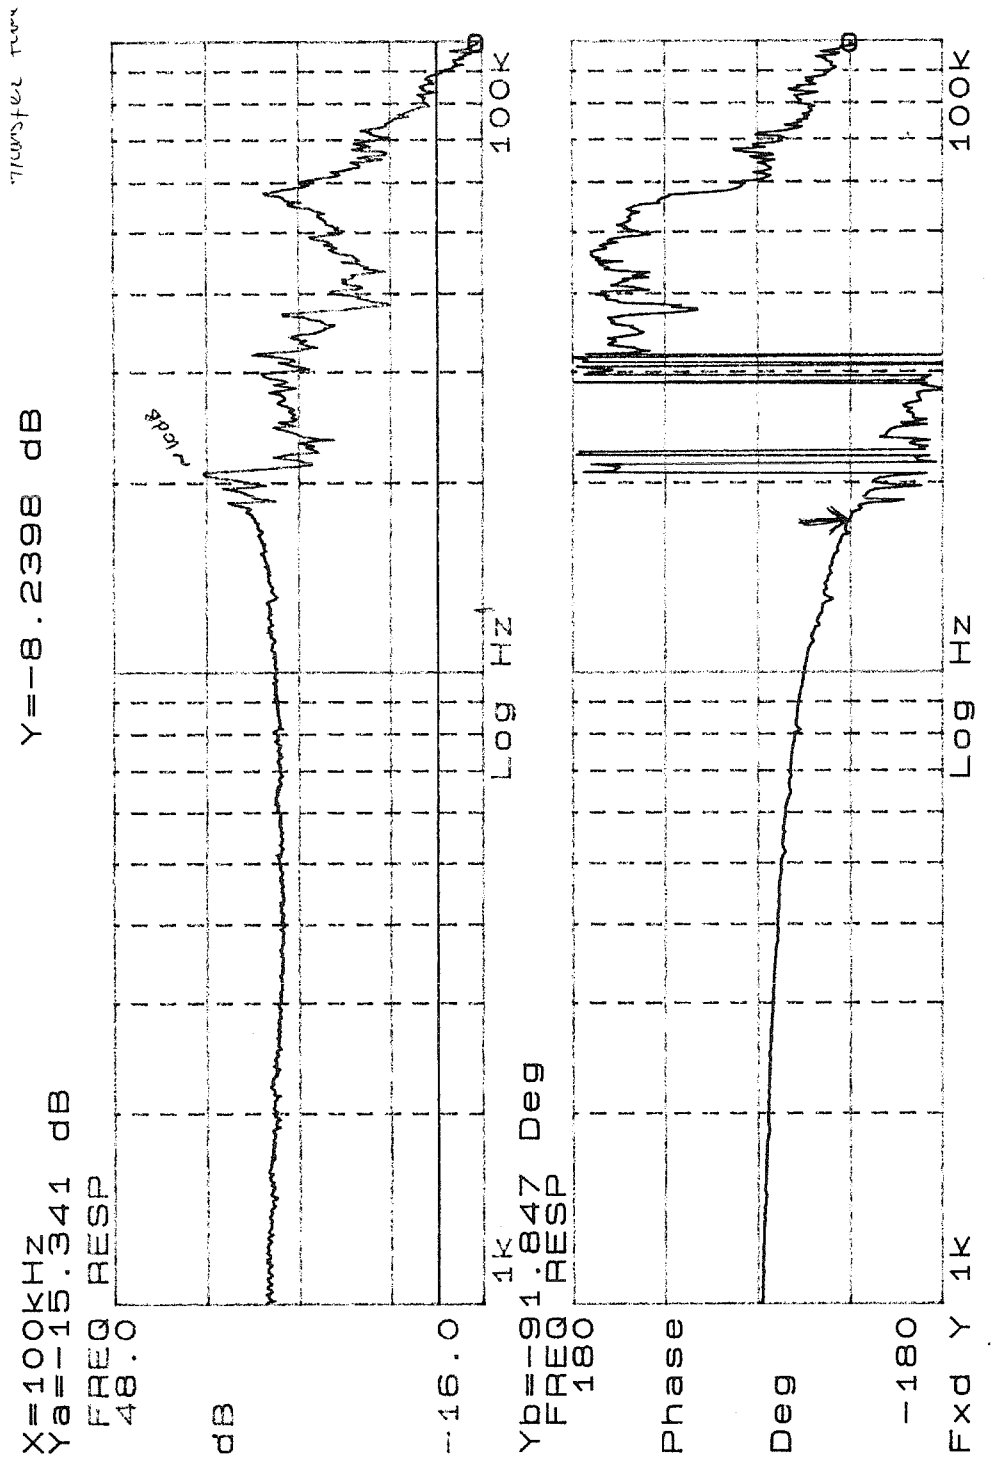

Figure 3.5: Open-loop response of the tweeter PZT in the Coherent 899-21 Ti:Sapph.

In order to be more aggressive at low frequencies, I've been bypassing the Coherent control box and now use a home-built servo for the tweeter. The Coherent servo never gets any steeper than 9 dB/octave, presumably in order to ensure that the loop would remain stable even when the Ti:Sapph was out of tune (implying reduced error-signal size). One potential drawback of bypassing the Coherent box is that you leave the galvo out of the game, but I've found that the tweeter has more than enough range to deal with DC drifts for hours at a time.

Figure 3.6 shows a schematic of the first servo amp I built for the tweeter, before we started using an external AOM as well. It's just a basic PI controller with a phase lead compensator to provide some phase margin at 6 kHz. My experience has been that it is generally better to implement separate proportional and integral branches and sum them, as opposed to using an "inline" design such as (integrator $\rightarrow$ low-pass $\rightarrow$ phase-lead network $\rightarrow$ output buffer). The latter design should theoretically give you exactly the transfer function you want, but the ones that I've built tend to misbehave (*e.g.*, differentiate) at frequencies above 20-30 kHz. My understanding of this is that it's just hard to find an op-amp that can really integrate at high frequencies, so you're better off not sending *all* of your feedback signal through the integrator. At some point I put together a test circuit for checking out performance of various op-amps as integrators, with a 300 kΩ input resistor and a 1 μF feedback capacitor (plus an output buffer). The winners seemed to be LT1122 and the good-old OP27, while LF356 and AD711 were not so impressive.

The overall design of Figure 3.6 uses two RC poles at 3 Hz to suppress the AC gain of the integral branch, so the high-frequency servo response is determined by the proportional branch alone. The proportional branch has two poles at 30 Hz to set the initial rolloff, plus a zero at 4 kHz to provide some phase margin and another pole at 8 kHz to set the ultimate slope at 12 dB/octave. The relative gains of the P- and I-branches are such that they cross around 30 Hz. This turns out to be perfectly stable, although an earlier version with the PI-corner at 12 Hz displayed monstrous low-frequency oscillations. Although I can't really claim to understand why, it seems that you have to keep the relative phase of the P- and I-branches *less than* 180 degrees

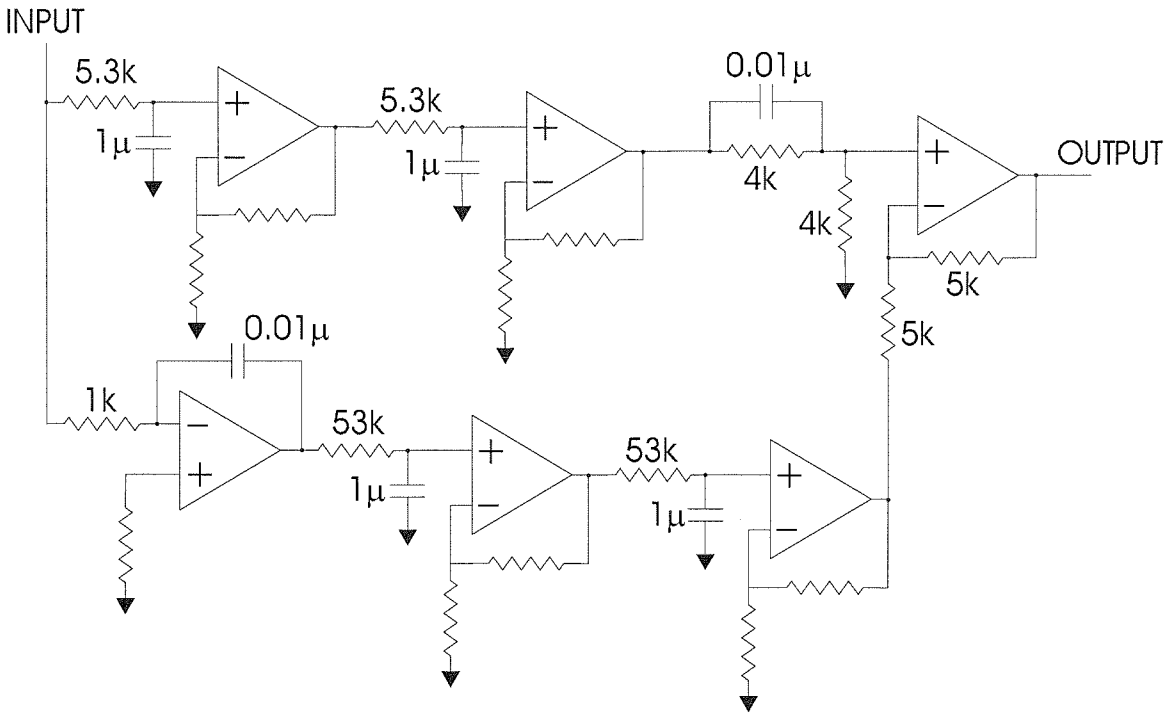

Figure 3.6: Schematic for the tweeter servo amplifier. All op-amps are OP27.

at the crossover point. For any particular design, it's worth checking out where the closed-loop poles will be before you start soldering—the Matlab Controls toolbox makes this a trivial exercise.

I found that it's just barely possible to lock the Ti:Sapph to the transfer cavity using the tweeter servo alone. As mentioned above, however, you have to unplug the brewster-plate galvo in order to do this, and even then the lock isn't terribly robust. When Jun arrived last year, we set up an external AOM to get faster control over the Ti:Sapph frequency. This lets you *really* crunch down the jitter. The Ti:Sapph light now gets double-passed through an 80-MHz AOM right after it comes out of the box. Using 15 cm lenses, and without terribly much attention to the alignment, we seem to get adequate pointing stability for even the probe beam that travels several meters on its way to the physics cavity. The AOM is driven by a Mini-Circuits VCO, which has acceptably low drift but not quite the modulation bandwidth that one would like (ours has -3 dB at about 70 kHz). Jun built a "bandpass" servo amp with a zero around 250 Hz (AC feedback coupling prevents DC pointing errors) and one

pole at about 2 kHz. The high-frequency rolloff of the servo loop ends up being much harder than 6 dB/octave, however, because of the sluggish VCO and because of the transfer cavity's 100 kHz linewidth. We generally see unity-gain oscillations in the neighborhood of 70-90 kHz (depending on the gain setting), and it's clear that having the AOM servo on lets us increase the gain of the tweeter servo without ringing the PZT resonances at 17 kHz. I'm not really sure what the tweeter-AOM crossover point is, but there don't seem to be any bad interactions between the two servo loops. With everything all locked up, the residual fluctuations of the Ti:Sapph-to-transfer cavity error signal suggest an rms linewidth less than 10 kHz. This number is supported by direct comparison of the Ti:Sapph to out-of-loop frequency discriminators such as the local oscillator cleaning cavity (FWHM=1 MHz).

Figure 3.7 shows the closed-loop spectrum of residual noise on the Ti:Sapph-to-transfer cavity error signal. Unity-gain oscillations from the AOM servo can be seen around 70 kHz, and the resonances of the Ti:Sapph tweeter PZT can also be seen around 17-20 kHz.

It's worth noting that a double-passed AOM servo really cuts into the available laser power, but we had no difficulty running the experiment with a total of $\sim$ 50 mW of stabilized light.

3.1.6 Photodetection

Up until the last few months of the experiment, we were using the same pair of photodetectors that Quentin used for all the work described in his thesis. The overall design, which originated sometime in the ancient history of Kimble group, is wonderfully simple and performs fairly well. I refer the reader to Quentin's thesis for the circuit diagram and such, but the detector basically consists of an FND-100 photodiode connected directly to a GPD-201 50Ω monolothic rf amplifier. One of the best features of this design is that excellent matching is automatically obtained between the gain and phase characteristics of any pair of such detectors. My experience has been that you can easily achieve $\geq$ 20 dB balancing over the heterodyne frequency

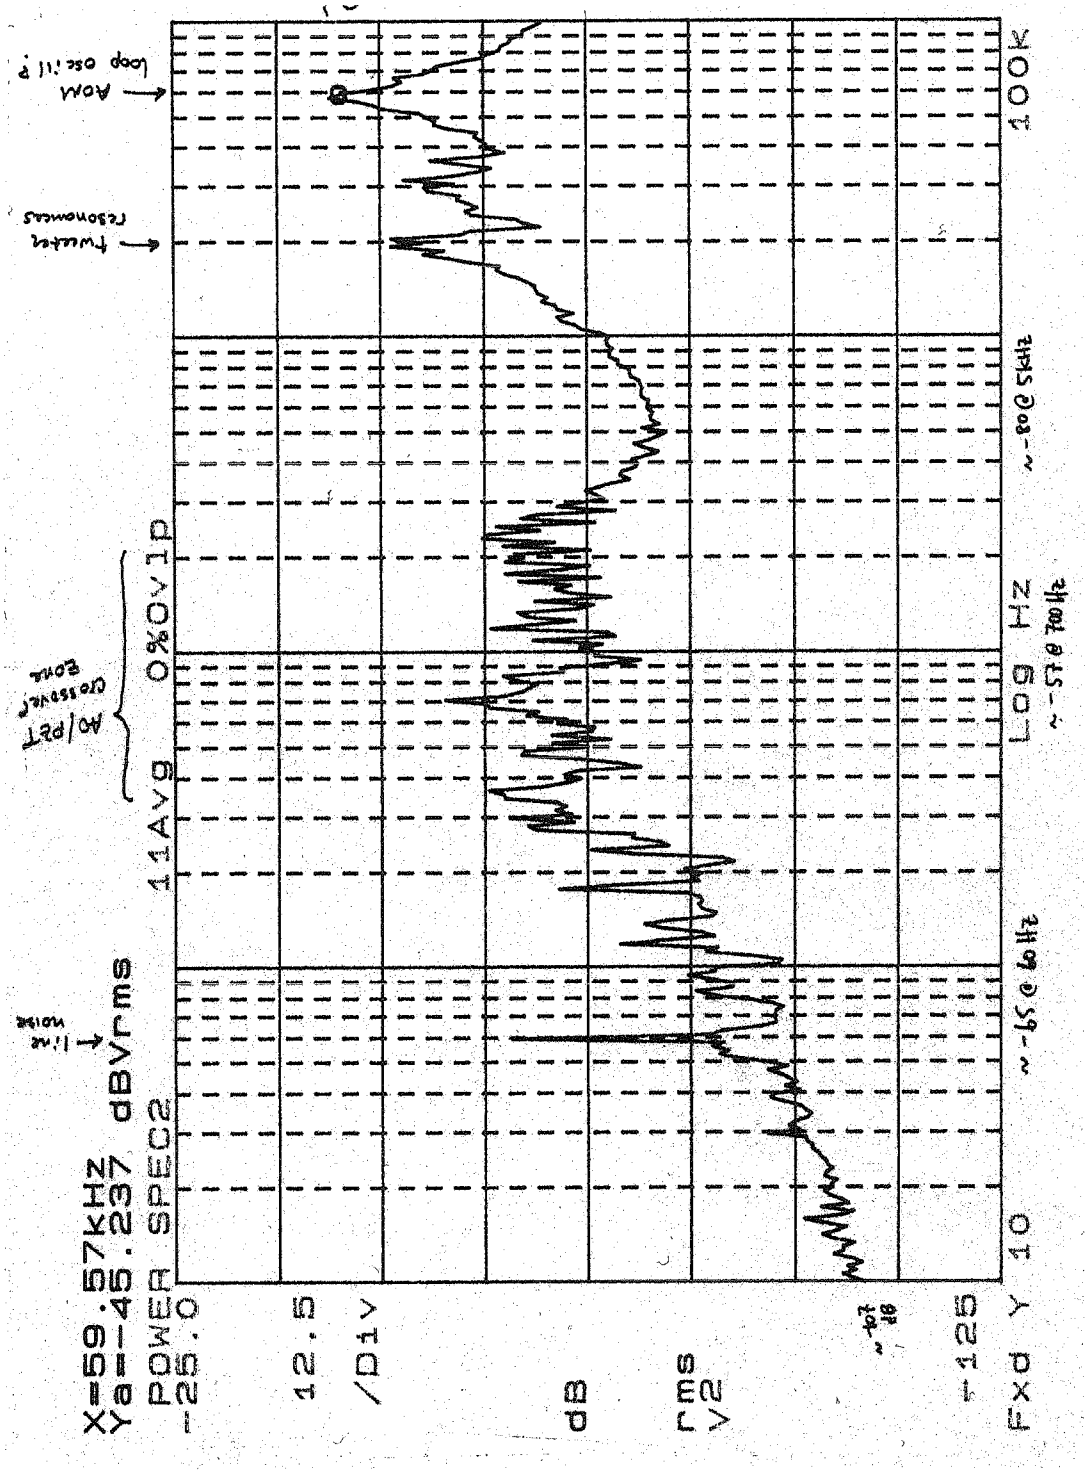

Figure 3.7: In-loop noise spectrum of the Ti:S locked to the transfer cavity.

range of $100 - 200$ MHz with detectors of this traditional design. On the down side, it's difficult to obtain more than 3 or 4 dB of shot noise with these detectors at high end of that frequency range. We typically used 4 mW of LO per detector in order to get as far above the electronic noise as possible, but this was really too much optical power.

If your goal is getting to the real shot-noise limit for recovery of the quadrature amplitudes, you'd really like the electronic noise to be down by 10 dB or more. As this clearly wasn't going to happen with the traditional 50Ω system, Jun recently built a new pair of heterodyne detectors based on the Signetics NE5211 transimpedance amplifier ($28k\Omega$ transimpedance gain). The photodiode is now a low-capacitance Hammamatsu S5982, and the detectors Jun built can see shot noise from about 5-300 MHz with 1 mW of LO. In the neighborhood of 50 MHz, it's possible to work 10 dB above the electronic noise. From $\sim$ 100 MHz on out though, the limit still seems to be about 5 dB. We measured the quantum efficiencies of the two Hammamatsu photodiodes we're using to be 0.68 and 0.71, although according to spec they ought to be more like 0.78 (the FND-100's were both about 0.86). Unfortunately, this pair of detectors turn out to have fairly different high-frequency rolloffs, both in terms of gain and of phase. This makes it difficult to achieve good balancing over broad ranges of frequency. I've certainly found that you can get 20 dB or so over a bandwidth of 10 MHz, as measured by the cancellation of AM written on the LO by joint action of an EOM and the local oscillator cleaning cavity.

We discovered early on that one has to be fairly careful not to saturate the AC gain of the transimpedence amps with technical noise on the LO. The frequency-servo AOM generates a fair amount of residual modulation on the LO at 76 MHz, for example. We actually have to strip this off in order to keep the photodetectors happy—hence the 1 MHz FWHM LO cleaning cavity. It's also necessary to be careful about the FM sidebands used for locking the cleaning cavity. We lock in reflection with weak FM sidebands at 24 MHz, which keeps the residual modulation of the transmitted beam acceptably low.

As noted in the previous chapter, it doesn't seem that we could ever get a very

good spatial overlap between signal beam and local oscillator. The best we did was a DC fringe visibility of 0.7, which is not quite up to Kimble Group standards from the squeezing era. All I can do at this point is offer some speculation as to why we did so badly. First, we propagated a fairly large (in the gaussian waist sense) beam from the physics cavity output to the heterodyne detectors. Due to geometric constraints, we had to allow the cavity output to expand for 30 cm before collimating it. Without really thinking about implications for the beam quality, we kept the beam like that and matched the local oscillator to it. The signal beam hit many (14) surfaces on its way to the detectors, and its large spatial extent meant that the signal beam wavefront was sampling abberations of these surfaces over a pretty large area. It probably would have been better to refocus the signal beam. A second possible problem was that our local oscillator cleaning cavity had a transverse mode spacing that was not so different from 1/3 its free spectral range. This means that even when we locked to a TEM_{00} error signal, there was probably a fair amount of transverse spatial mode leaking through to compromise the overlap with the signal beam, which one would like to think was a good TEM_{00} when it left the physics cavity.

For anyone who might be interested, the rf components that we used for processing the photocurrent are as follows. The outputs of the two photodiodes first went through identical rf turret-attenuators (HP) to balance the gains, with different lengths of cable used to match the phases as well. The signals were then subtracted in a 180° splitter (Mini-Circuits ZFSCJ-2-1), and the difference photocurrent was boosted by an rf amplifier (ZFL-500LN). The amplified current then goes to a 0° splitter (ZFSC-2-1), whose two outputs are used to generate the pair of orthogonal quadrature amplitudes. We use two double-balanced mixers (ZLW-1H) to bring the heterodyne beat note down to baseband, where the two rf local oscillators are produced by running the maximum output power from an HP8656B signal generator through a 90° splitter (ZMSCQ-2-180, -120, or -90 depending on the frequency). The two quadrature amplitude signals then go to separate SR560 preamplifiers, which are used to bring them up to ± 0.5 V.

We chose these particular components so as to have a completely SMA-connected

system (we had problems with flakey ground contacts in an earlier BNC version), to have the 0° and 180° splitters be as broadband as possible, to minimize insertion losses, and to have mixers that went all the way down to DC for signals coming out the IF ports. Note that the ZLW-1H's are level-17 mixers, but we could only give them around +13.5 dBm of rf local oscillator. The phase accuracy of the 90° splitters wasn't so great, and they were relatively narrowband. So these would be the components to upgrade in the future. It looks like a company called Merrimack could supply much more sophisticated devices for the job, but they would be expensive ($\sim$ \$6 – 7k) and had a lead time of several months.

3.1.7 The vacuum system

I originally ordered this vacuum chamber for an effort to see optical bistability using an atomic beam and quartz microspheres. It's a copy of the chamber that Zhen Hu and Mike Winters designed for cooling and trapping experiments, built out of 303 stainless steel by NorCal. The chamber is essentially a feedthrough collar, 4 inches in height with a 6-inch conflat flange on either end. Eight 2.75-inch conflat ports look radially into the center of the chamber, providing an optical axis every 45 degrees. Figure 3.8 shows the overall configuration of the vacuum system as we've used it for the past two years. Even with the physics cavity, vibration-isolation stack, and Cesium oven, a single 20 l/s Starcell ion pump suffices to keep the background pressure in the low 10^{-8} torr (as estimated from the pump current).

The original idea of the vacuum system design was to have a setup that could easily be switched between running an old-style thermal beam experiment and a dropping-MOT experiment, by raising or lowering the oven and collimating coldfinger relative to the chamber with the cavity in it. We never actually used this flexibility, however. It would be much nicer to have a differentially-pumped system with a Cs vapor cell to load the inital MOT, as much of the hassle of running the experiment from day to day was having to heat up and cool down the Cs oven, and having constantly to refill the dry ice box for cooling the coldfinger. Heating up the oven, and also just

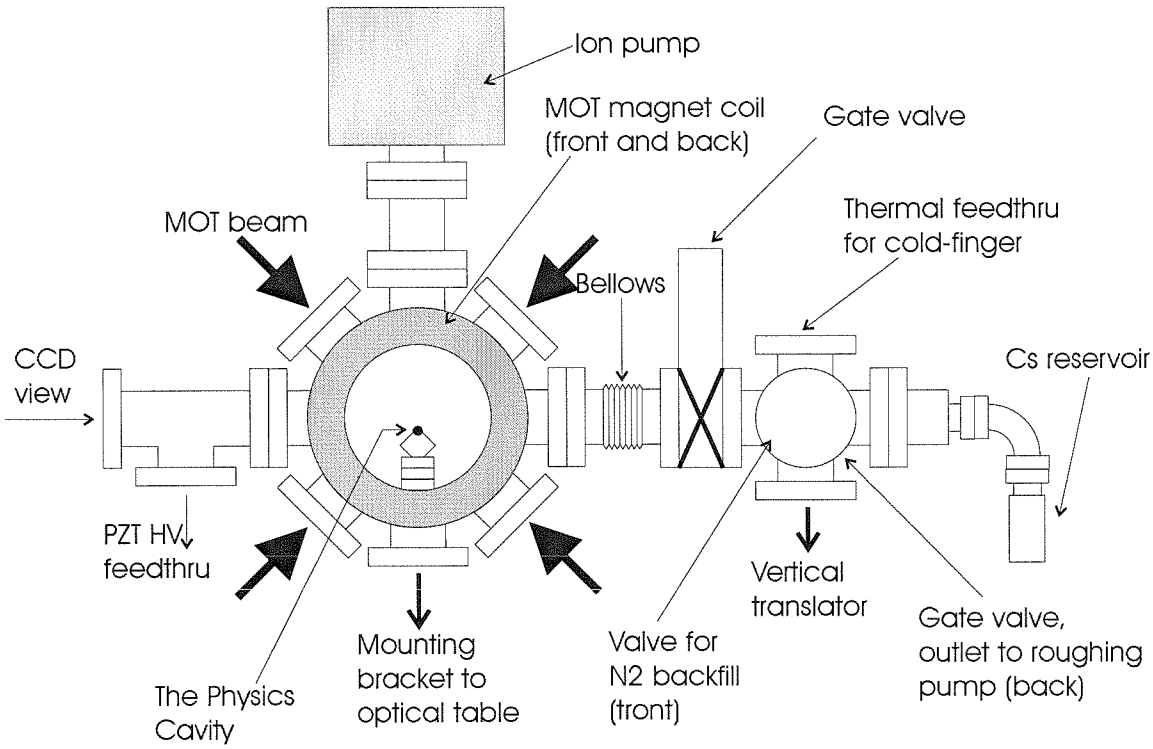

Figure 3.8: The vacuum system, drawing not to scale!

having all that extra surface area in the vacuum system, certainly also compromised the base pressure.

I think it was important to really optimze the vacuum conductance from the cavity chamber to the ion pump, as we couldn't seem to get as low a base pressure in the chamber with previous configurations of the hardware that had the pump in a different place. Having the gate valve between the cavity chamber and cesium oven was a great convenience—it allowed us either to change the cesium without bringing the chamber up to air, or to make adjustments to the cavity without oxidizing the cesium.

3.1.8 The trapping lasers and magnetic coils

The light for the magnetooptic trap is generated by a pair of grating-stabilized SDL 5411 diode lasers, each of which is mounted in a slightly-modified version of Ken Libbrecht's laser housings and driven by Libbrecht-Hall current controllers [62]. In

order to make the grating adjustments a bit more convenient, I've been using Lees Optical LM-1 mounts inside the laser housing in place of the original flexure mounts. I should also note that since the current controller was originally designed for positive-polarity diode lasers (whereas the SDL diodes are negative), you have to ground the diode and run the case positive in order to use the controllers without modification. With the old STC diode lasers we would always remove the "window" (which was really only meant to be a protective cap), which allowed us to use short working-distance Melles-Griot compound collimating lenses (06-GLC-001/B). The windows on the SDL diodes are not really removable, and are optically very nice anyway, so I've been leaving them on and using long working-distance aspheres that are available from Thor Labs (C230 TM-B). I used to worry that exposing the diode facets to air led to some danger of water condensation, so overall I like the present configuration much better. I think you can see some effects of optical reflections from the window when running the diodes at low current, but well above threshold they don't seem to be a problem.

One nice thing about mounting the diffraction grating (1800 grooves/mm, holographic/VIS from Edumnd Scientific) on an LM-1 is that you can scan the entire wavelength range of the diode just by turning the horizontal knob, without having to stop and redo the vertical alignment. As a consequence, you can always plug in a brand new diode and pull it to the Cesium resonances without ever using a Wavemeter. You simply optimize the threshold at some arbitrary grating angle, turn the current up to the desired setpoint, put a Cesium cell right at the laser output, tune the grating angle until you see a flash in the cell (with an IR viewer), then re-optimize the threshold at the correct grating angle. If you scan the grating and tweak the current you should immediately be able to find lines in a saturated-absorption setup.

I used one diode to produce trapping light on the Cesium $6S_{1/2}, F = 4 \rightarrow 6P_{3/2}, F' = 5$ transition at 852.359 nm, and the second to produce a repumping beam on the $F = 3 \rightarrow F' = 4$ transition at 852.339 nm. I locked the diode frequencies by feeding back to both the grating PZT (at something like 10-100 Hz) and the

diode current (at something like 10-100 kHz). The DC drifts were small enough that I didn't bother to use an integrator. I got an error signal for the repumping diode by modulating the current at 10 MHz and looking at saturated-absorption lines in FM. In order to have an FM error signal for the trapping laser as well, *without* putting sidebands on the trapping light (and without investing in an EOM), I made use of optical-pumping crossovers to lock the trapping diode frequency relative to that of the repumper.

The setup for doing this is shown in Figure 3.1. The repumping diode is first locked to the $3 \rightarrow 4$ transition frequency using a standard FM saturated-absorption setup. Some "probe" light from the repumping diode is then sent through a second Cs vapor cell to monitor the absorption by atoms with zero velocity along the z-axis. This type of signal should actually be quite weak, as atoms traversing the probe beam will quickly be pumped into the dark $F = 4$ ground hyperfine level. But we then also introduce an overlapping, counterpropagating "pump" beam into the cell, derived from the trapping laser. The probe absorption increases drastically if the "pump" beam is tuned to either the $4 \rightarrow 4$ or the $4 \rightarrow 3$ transition, due to optical repumping from the $F = 4$ hyperfine ground level back to the $F = 3$. Hence, FM detection of the transmitted *probe* generates an error signal for stabilization of the *pump* beam frequency. The absorption peaks have about 30 MHz FWHM when using beams of ~ 1 mW/cm^2 intensity, with much better contrast than saturated-absorption lines and negligible Doppler background.

To end up with light at a frequency appropriate for trapping, I send a 140-MHz downshifted beam from an AOM into the locking setup. This means that the trapping laser itself sits at the $4 \rightarrow 4$ transition plus 140 MHz, and the light for the MOT is generated as the first-order upshifted beam from a 100 MHz AOM (resulting in a 10 MHz red-detuning from the $4 \rightarrow 5$ transition). The latter AOM also serves as a fast shutter for turning the MOT on and off, and I have found that a Mini-Circuits rf switch (ZYSWA-2-50DR) has sufficient attenuation (around 45 dB at 100 MHz rf) to turn the light off so that there's no residual optical molasses in the chamber—the Isomet drivers typically don't. The Isomet AOM's suffer from thermal problems that

cause the pointing of both the diffracted and zeroth-order beams to depend on the average dissipated rf power. Hence, if you align the MOT beams with one setting of the rf duty cycle (*e.g.*, 100%), you can't change the duty cycle without messing things up. I've heard that AOM's from NEC or Crystal Technology are much better in this regard.

We don't actually want to start with too many atoms in the trap for the purposes of doing cavity QED, so I have generally used 3-4 mW in ~ 1 cm^2 trapping beams during alignment and ~ 10 μW/cm^2 beams while running the experiment. The MOT loads directly from an uncooled atomic beam, which effuses from a Cesium oven at ~ 70 C through a 200 μm pinhole. The atomic beam is collimated by a pair of 3 mm aperatures that we keep at dry-ice temperature. We typically load the MOT for 1 sec each data cycle. My experience with this MOT configuration has been that you get plenty of atoms ($\sim 10^5$) as long as the background pressure of the chamber remains below 2×10^{-7} torr, and we typically had pressures around $\sim 4 \times 10^{-8}$ torr.

In order to actually drop the MOT—as opposed to accidently launching it in some random direction—I found that the one crucial point is to make sure the trapping light is completely off before changing the magnetic fields. We typically used an axial field gradient of 30 g/cm, generated by a pair of coils (~ 7 inch diameter, originally designed by Wolfgang Lange) in the usual anti-Helmholtz configuration (6 amps into a total of 250 windings). In order to avoid excessive transients when switching off the current, we put a capacitor on the programming input of the current supply that combined with the input resistance to yield a ~ 4 ms RC time constant. Even with that precaution, we at times had trouble with *mechanical* vibrations associated with the magnet switching. Our magnetic coils were designed simply to hang on the vacuum chamber, without any additional mechanical support structure. When the current switches on or off, the two coils pull or push each other and the mechanical impulse gets transmitted to the vacuum chamber and also the physics cavity itself. In the early days this would actually unlock the cavity every once in a while. With the improved servos this no longer occurs, but you can clearly still see transients in the physics cavity error signal. On rare occasions a faint tail of this mechanical ringdown

would actually persist into the data window, but the residual oscillations were very small and of low frequency ($\sim$ 300 Hz).

So mounting magnetic coils directly on the vacuum chamber is a bad idea. In addition to unwanted mechanical impulses, we sometimes had problems with the coils getting hot. This would then heat up the chamber itself, leading to noticeably-increased outgassing and eventually to thermal drift of the physics cavity length. Wolfgang actually designed these coils to be water-cooled, but we generally found that running water through the coilforms also leads to significant vibration of the chamber. Above and beyond these sorts of mechanical problems, the primary drawback of the way we mounted the magnetic coils was that the null of the anti-Helmholtz field became fixed with respect to the geometry of the chamber. It just wasn't practical to try to install the physics cavity such that the gap between the mirrors was really lined up with the point at which the MOT would naturally form, so in the end we had to align the MOT to the cavity by applying small bias magnetic fields. Physically moving the anti-Helmholtz coils themselves would have been a much better way to do this alignment, but again this just wasn't practical given their extreme weight and the intricate tangle of optic mounts that had to press right up against the windows of the vacuum chamber. So our ability to trim the net magnetic field at the physics cavity was compromised by the need to put the MOT in the right place, and we ended up with a Zeeman quantization axis that pointed $\sim$ 30 degrees off of the cavity axis. This wasn't really such a tragedy in our case because the cavity was birefringent anyway, but in the future this sort of problem has to be anticipated and avoided.

Back while Quentin was still here, we briefly investigated the use of large homogeneous magnetic fields to produce frequency splittings between different Zeeman transitions. Although the huge magnetic coils we have are clearly overkill for the simple purpose of making a MOT, they make it easy to produce $\sim$ 60 gauss homogeneous fields at the physics cavity. Taking advantage of our bipolar current supply, we connected the coils in series with a full-wave rectifier such that they produce an anti-Helmholtz field for positive applied voltages, and a Helmholtz field for negative applied voltages. The circulation of current in one of the two coils changes with the

sign of the applied potential, but stays the same in the other by action of the rectifier. We didn't actually use this trick for much in the end, but I thought it was worth passing along.

3.1.9 Data acquisition

In our initial efforts, we used a National Instruments multi-function I/O card to digitize and record the data. At the time, NI offered two versions of the board—one with 250 kHz sampling rate and another with 500 kHz (both with 12-bit resolution). I somehow felt that it was nice but extravagant to have bought the faster model, but once the experiment actually started working it quickly became clear that we were severely undersampled even at 500 kHz.

For a while we used an 8-bit ADC sold by Keithley-Metrabyte, which has two independent channels that can simultaneously sample at up to 100 MHz. Although we learned some things by looking at higher-bandwidth data with that card, eight bits just doesn't provide enough dynamic range. For data with gaussian noise, you basically have to sacrifice one bit to prevent the signal from clipping, and you're then left with a resolution of only one percent.

We presently use a 12-bit ADC from Gage Applied Science (model 8012/A), which also has two independent channels that can simultaneously sample at up to 100 MHz. The card has 1 Mb of onboard RAM per channel, so it's no problem to take 50 ms data windows at 10 megasamples per second per channel. We were fairly conservative in setting the analog signal bandwidth, in order to be completely sure that we were not aliasing. In addition to using a 1.9-MHz single-pole filter on the input of the SRS560 voltage preamplifier for each quadrature amplitude signal, we set the preamp's built-in filters to roll off at 12 dB/octave with a 300 kHz corner. At the Nyquist frequency of 5 MHz, we thus had $6 + 12 \times \log_2(5/0.3) \approx 54.7$ dB attenuation from the SRS560 plus an additional $3 + 6 \times \log_2(5/1.9) \approx 11.4$ dB from the rf filter on the input. Fourier components of the signal at frequencies above Nyquist were thus attenuated by more than $10^{-66.1/20} \approx 5 \times 10^{-4}$. Given that the SR560 gains were set to keep

the rms signals (which contain a large DC component) well within the dynamic range of the ADC, this ought to provide plenty of anti-aliasing relative to the 12-bit ADC resolution ($\sim 2 \times 10^{-4}$).

Before looking at the quadrature amplitudes data, it is necessary to numerically remove drifts in the phase of the rf local oscillator relative to the heterodyne signal. The basic strategy I've been using is to say that on timescales that are slow compared to atomic transits but fast compared to any phase drifts, the phase quadrature of the heterodyne photocurrent should be forced to have a mean value of zero. Hence the following Matlab algorithm. First, read in the two digitized signals $y1, y2$, which always represent an orthogonal pair of quadrature amplitudes but have an arbitrary, time-varying phase ϕ relative to the actual amplitude and phase quadratures. Next, compute ϕ by taking the arctangent of $-y2/y1$, as *unwrap(atan2(y1,−y2))*. Then apply a running boxcar average (2 ms width) to smooth out $\phi(t)$. Finally, compute "flattened" versions $x1, x2$ of the raw signals $y1, y2$ by applying a time-varying rotation,

$$\begin{pmatrix} x1 \\ x2 \end{pmatrix} = \begin{pmatrix} \cos\phi & -\sin\phi \\ \sin\phi & \cos\phi \end{pmatrix} \begin{pmatrix} y1 \\ y2 \end{pmatrix}. \tag{3.2}$$

Chapter 4 Studies in quantum measurement theory

4.1 Broadband measurements

Abstract: In this section, I utilize the Caves-Milburn model for continuous position measurements to formulate a broadband version of the Standard Quantum Limit (SQL) for monitoring the position of a free mass, and illustrate the use of Kalman filtering to recover the SQL for estimating a weak classical force that acts on a quantum-mechanical test particle under continuous observation. These derivations are intended to clarify the interpretation of SQL's in the context of broadband quantum measurement, with particular attention paid to the question of how it might be possible to *verify* that a given laboratory measurement does indeed achieve backaction-limited sensitivity. The method used to analyze force detection may be extended to the case of weak classical forces with arbitrary time dependence, including nonstationary and impulsive signals.

Substantial efforts have been devoted to elaborating Standard Quantum Limits (SQL's) for both discrete [63, 28, 64, 5] and continuous [5, 65] measurement of the position of a free mass. In large part, the motivation for such investigations stems from a pressing need to identify any possible constraints imposed by the principles of quantum measurement on the experimental possibility to detect gravitational waves [66, 67]. My objectives in this paper will be to formulate an SQL appropriate to *broadband* continuous measurements of the position of a free mass (expressed in terms of the signal bandwidth and with measurement sensitivity given in units of length per $\sqrt{\text{Hz}}$), to discuss its proper intepretation, and to demonstrate its compatibility with the usual SQL [5] for detecting a weak classical force. The mathematical analysis will be adapted to the continuous measurement model of Caves and Milburn [38], which

appears to have direct relevance for concrete experimental scenarios such as atomic force microscopy [68] and cavity QED [16, 69].

Although the limits I derive will be familiar from previous studies of the discrete and narrowband measurement scenarios, my emphasis here will be on formulating these limits in a manner that is specific to broadband quantum measurement. In particular, I have found that some care needs to be taken in deriving an SQL for detecting weak forces from the SQL for continuous position measurement. The method I use below emulates the classical technique of Kalman filtering [70, 71], and was motivated by the general strategy of quantum system identification discussed in [14]. Throughout this paper I have adopted the time-domain, state-space perspective of quantum measurement whose virtues have become apparent from recent advances in adaptive quantum measurement [15] and quantum feedback [40]. Such methods hold great promise for the challenge of formulating experimentally tenable strategies for broadband quantum nondemolition measurement.

The mathematical formalism necessary to treat continuous quantum measurements has been developed by numerous authors, with the most relevant works for the present discussion being [38, 71, 72]. While it is not absolutely essential to go to the continuous limit, doing so will allow us to use the convenient notation of stochastic differential equations (SDE's) [73]. What really matters for the discussion at hand is an assumption that the timescale associated with measurement interactions and readouts is much shorter than any timescale on which we wish to understand the system dynamics.

In the Caves-Milburn model for continuous quantum-mechanical measurements of position [38], there is one system of interest, and an infinite succession of identical (and identically-prepared) "meters." Let the meters be labelled by an index r. The system is brought into momentary contact with the r^{th} meter at time $t_r = r\tau$, and the position operator of each meter is measured sharply just after it has interacted with the system. The string of measurement results thus generated constitutes a classical record of the system evolution. Continuous measurement is achieved in the limit where the time-interval between measurements τ goes to zero, with the uncertainty

of each individual measurement simultaneously going to infinity in an appropriate manner.

Let x, p denote the system position and momentum operators, and X_r, P_r the position and momentum operators of the r^{th} meter. We adopt the following time-dependent interaction Hamiltonian between the system and the meters:

$$H_{int} = \sum_{r=1}^{n} \delta(t - r\tau)\, xP_r. \tag{4.1}$$

The delta-function form is chosen to facilitate the limit $\tau \to 0$, but as discussed above it may be viewed as the idealization of any "shrinkable" function with compact support. We assume that each meter, just before it interacts with the system, is prepared in the pure state $|\Upsilon_r\rangle$ with Gaussian wave-function ($X_r |\xi_r\rangle \equiv \xi_r |\xi_r\rangle$)

$$\langle \xi_r \mid \Upsilon_r \rangle \equiv \Upsilon(\xi_r) = \frac{1}{(\pi\sigma)^{1/4}} \exp\left[\frac{-\xi_r^2}{2\sigma}\right]. \tag{4.2}$$

Caves and Milburn have have derived an exact expression for the conditional evolution of the system state under the following measurement protocol: at each time t_r, *i.* couple the system to the r^{th} meter prepared in state (4.2), *ii.* evolve the system and meter under the interaction Hamiltonian (4.1), *iii.* perform a precise measurement of X_r. If we write $\rho(t_r-)$ for the system state just before the measurement at time $t_r \equiv r\tau$, the post-measurement system state $\rho(t_r+)$ is given by

$$\rho(t_r+) = \hat{\Upsilon}(\xi_r)\, \rho(t_r-)\, \hat{\Upsilon}^\dagger(\xi_r), \tag{4.3}$$

$$\hat{\Upsilon}(\xi_r) \equiv \langle \xi_r \mid e^{-ixP_r/\hbar} \mid \Upsilon_r \rangle, \tag{4.4}$$

where ξ_r corresponds to the actual outcome of the meter projection in stage *iii.* Recall that x, X_r, P_r are Hermitian operators, ξ_r is a c-number, and $|\Upsilon_r\rangle$ is a Hilbert-space ket. $\hat{\Upsilon}_r$ is thus a quantum operation density. Qualitatively speaking, the basic effect of (4.3) is to *shift* the centroid $x'_r \equiv \mathrm{Tr}[x\rho(t_r-)]$ of the system's position-space distribution *towards* the value conveyed by ξ_r, and to reduce its overall width

$\Delta_r' \equiv \mathrm{Tr}\left[x^2\rho\left(t_r-\right)\right] - \left(x_r'\right)^2$ by an amount that depends on the ratio Δ_r'/σ. Note that this is not a projective measurement—the post-measurement state $\rho\left(t_r+\right)$ depends on both the measurement result ξ_r *and* the pre-measurement state $\rho\left(t_r-\right)$.

The operation (4.4) maps Gaussian pure states of the system to Gaussian pure states, so we can in fact parametrize the selective evolution of a free particle (initially prepared in a Gaussian pure state) by just four real numbers. Chosing $x_r \equiv \langle x\rangle, p_r \equiv \langle p\rangle, \Delta_r \equiv 2\left\langle(\Delta x)^2\right\rangle$, and $\varepsilon_r \equiv \langle \Delta x\Delta p - \Delta p\Delta x\rangle/\hbar$, Caves and Milburn derived a set of difference equations for the evolution of these quantities with each successive measurement. They further showed that Gaussian initial states generally converge toward stationary evolutions, in which x_r and p_r evolve stochastically but Δ_r and ε_r vary periodically—during the time intervals between measurements both widths increase according to the free Hamiltonian, but the effect of each measurement is to reduce them by a constant "contraction factor" C such that $\Delta_r \equiv \mathrm{Tr}\left[x^2\rho\left(t_r+\right)\right] - \left(x_r\right)^2$ is independent of r. The form of the evolution equations is especially convenient in this stationary regime, so in what follows we shall assume that stationarity has been bootstrapped by preparing the system in an appropriate initial state.

In order to derive a broadband SQL for position measurements, we start with the stationary difference equations (3.30a,b) from [38]:

$$\begin{aligned} x_{r+1} - x_r &= p_r\frac{\tau}{m} + \frac{C-1}{C}\left(\xi_r - x_r'\right), \\ p_{r+1} - p_r &= \frac{\hbar}{\sigma\sqrt{C}}\left(\xi_r - x_r'\right). \end{aligned} \tag{4.5}$$

Here unprimed quantities x_r, p_r describe the system state just after the $(r-1)^{th}$ measurement has been made, and x_r' holds just before the r^{th} measurement.

We first need to rewrite the difference equations (4.5) as a stochastic differential equation (SDE), by taking $\tau \to 0$ together with $\sigma \to \infty$ such that their product $D \equiv \sigma\tau$ stays constant. Recall, from equations (3.20,3.21) of [38], that $\left(\xi_r - x_r'\right)$ is a Gaussian random variable with zero mean and variance $\sigma C/2$. In the continuous limit it also turns out that $C \to 1$ [38], so we immediately have an explicit (Ito) stochastic

differential equation:

$$d\begin{pmatrix} x \\ p \end{pmatrix} = \begin{pmatrix} p/m \\ 0 \end{pmatrix} dt + \begin{pmatrix} 0 \\ \sqrt{\frac{\hbar^2}{2D}} \end{pmatrix} dW, \tag{4.6}$$

where dW is a Wiener increment [73].

To derive an SDE for the rms position, we begin by applying Ito's formula to derive an SDE for $x_{ms} \equiv \langle x^2 \rangle$,

$$\begin{aligned} d\left(x^2\right) &= \frac{2p}{m} x\, dt, \\ \frac{d}{dt} x_{ms} &= \frac{2}{m} \langle px \rangle . \end{aligned} \tag{4.7}$$

Substituting x, p by integrals of equation (4.6), and keeping in mind that the "functions" $dW(t)$ appearing in the expressions for x and p will be identical for any given stochastic realization,

$$\begin{aligned} \frac{d}{dt} x_{ms} &= \frac{\hbar^2}{m^2 D} \left\langle \int_0^t dW(t') \int_0^t dt'' \int_0^{t''} dW(t''') \right\rangle \\ &= \frac{\hbar^2}{2m^2 D} t^2. \end{aligned} \tag{4.8}$$

Hence,

$$x_{rms}(t) = \frac{\hbar}{m\sqrt{6D}} t^{3/2}. \tag{4.9}$$

We now turn to quantify the fundamental measurement noise inherent in the present scheme. As mentioned above, the discrete measurement errors (before taking the continuous limit) are a stationary Gaussian process with variance $C\sigma/2$. Hence, in any given time interval B^{-1} the sample variance for the $N \equiv (B\tau)^{-1}$ measurements will be

$$\begin{aligned} \left\langle \sum \xi_r^2 \right\rangle &= \frac{C\sigma}{2B\tau}, \\ \Xi_{rms} &\equiv \frac{1}{N} \xi_{rms} = \sqrt{\frac{C\sigma B\tau}{2}} \to \sqrt{\frac{DB}{2}}. \end{aligned} \tag{4.10}$$

Note that $\sqrt{D}$ thus represents a measurement inaccuracy (or "sensitivity," depending on your point of view) in units such as meters/$\sqrt{\text{Hz}}$. In the sequel we shall regard $\Xi(t)$ as the time-domain measurement signal, so that B should be interpreted as a low-pass *bandwidth*. It is important to understand that D and B are completely independent variables—D parametrizes the system-meter coupling strength (a physical quantity), and B, the degree of smoothing applied to the measurement results (a signal-processing quantity).

We can now formulate a condition for backaction-induced diffusion to become visible against the fundamental measurement noise:

$$\begin{aligned} x_{rms} &= \Xi_{rms} \Rightarrow \\ \frac{\hbar}{m\sqrt{6D}}\, t_*^{3/2} &= \sqrt{DB/2}, \\ t_* &= \left[\sqrt{3B}\frac{Dm}{\hbar}\right]^{2/3}. \end{aligned} \tag{4.11}$$

This expression should be taken to predict the overall measurement time t_* at which backaction-diffusion should cause the rms wander x_{rms} of the observed particle (with mass m) to equal the point-to-point scatter Ξ_{rms} in continuous measurements made with measurement inaccuracy $\sqrt{D}$ and filtered to bandwidth B. Viewed as such, I claim that equation (4.11) represents the most sensible form for a *broadband* SQL for continuous observation of the position of a free mass. Note that the timescale for backaction noise to manifest itself increases with increasing signal bandwidth (measurement noise hides the backaction), increases with particle mass (scenario is less "quantum"), and increases with measurement inaccuracy (need a good measurement to get to the SQL). This picture is illustrated by Figure 4.1a (for fixed D, B, and m), which shows $x_{rms}(t)$, Ξ_{rms}, and t_* together with two stochastic realizations of $x(t)$.

In the special case where we choose to integrate the measurement results over all of a given observation interval t, $B \equiv 1/t$ and we recover a single measurement scenario. We can then derive a basic figure-of-merit for continuous measurements—defining $S \equiv \sqrt{D}$ as the measurement inaccuracy and keeping t as the total observation time,

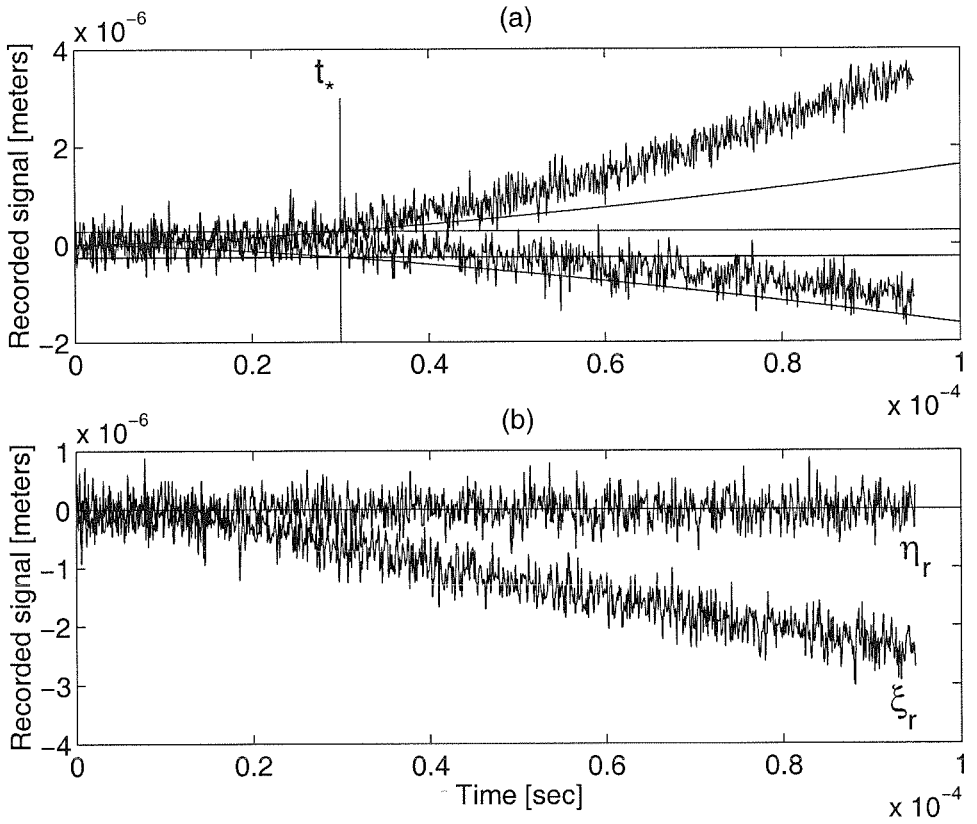

Figure 4.1: (a) Horizontal lines indicate the rms signal noise $\pm\sqrt{(DB/2)}$, curved lines indicate $\pm x_{rms}(t)$, and the vertical line indicates t_* for $D = 1.42 \times 10^{-20}$, $B = 10^7$, and $m = 2.22 \times 10^{-25}$ (mks units). Two stochastic realizations of the measurement signal are also shown, generated by direct simulation of equations (19)-(21). (b) The nonstationary component of the signal ξ_r can be perfectly subtracted away to yield $eta_r \equiv \xi_r - \hat{x}'_r$, which is a Gaussian noise process with variance $\sigma C/2$.

we find that we need

$$\frac{S}{t} \leq 3^{-1/4}\sqrt{\frac{\hbar}{m}} \tag{4.12}$$

in order for noise in the measurement record ultimately to be dominated by real displacements of the particle due to backaction diffusion. Alternatively, we can write down something like an inference-disturbance relation:

$$\Xi_{rms} \cdot x_{rms} \geq \sqrt{\frac{D}{2t}} \cdot \frac{\hbar}{m\sqrt{6D}}\, t^{3/2} = \frac{\hbar t}{2\sqrt{3}m}. \tag{4.13}$$

This relation asserts that the product of the integrated measurement accuracy and the rms displacement of the test particle due to backaction-diffusion must exceed a

lower bound that increases linearly with time.

We now turn to consider the estimation of weak forces acting on the test particle, based on something like Kalman filtering of the measurement record. We begin by writing down a modified form of the stationary difference equations (still explicit),

$$x_{r+1} - x_r = p_r \frac{\tau}{m} + \frac{C-1}{C}(\xi_r - x'_r), \tag{4.14}$$

$$p_{r+1} - p_r = \alpha\tau + \frac{\hbar}{\sigma\sqrt{C}}(\xi_r - x'_r). \tag{4.15}$$

Here we have simply added a uniform force field α that acts on the particle at all times [74]. Our task will be to detect $\alpha \neq 0$.

Based on the measurement results ξ_r, we want to update recursive estimators $\hat{x}_r$ and $\hat{p}_r$ for x_r and p_r, respectively. Assuming noiseless readout of ξ_r, the update equations read

$$\hat{x}_{r+1} - \hat{x}_r = \hat{p}_r \frac{\tau}{m} + \frac{C-1}{C}(\xi_r - \hat{x}'_r), \tag{4.16}$$

$$\hat{p}_{r+1} - \hat{p}_r = \frac{\hbar}{\sigma\sqrt{C}}(\xi_r - \hat{x}'_r). \tag{4.17}$$

Note that we do not include α in the update rules, as it is assumed to be an unknown quantity. In order to infer α from the recursive estimators $\hat{x}_r$ and $\hat{p}_r$, we must focus on the behavior of the quantity $\eta_r \equiv \xi_r - \hat{x}'_r$. As we know the distribution for ξ_r, we can write

$$\eta_r = (x'_r - \hat{x}'_r) + \sqrt{\frac{C\sigma}{2}}\mathcal{N}_r[0,1], \tag{4.18}$$

where $\mathcal{N}_r[0,1]$ is a Gaussian deviate with zero mean and unit variance. Hence, η_r will generally be the sum of a Gaussian deviate and an uncorrelated process (the "error-signal") that reflects the cumulative inaccuracy of our recursive position estimator.

Note that if there is no external force acting on the particle ($\alpha = 0$) there is no measurement inaccuracy, and the recursive estimator $\hat{x}_r$ can be used *perfectly* to subtract the backaction-induced diffusion from the signal ξ_r (see Figure 4.1b). Indeed, demonstrating the ability to do so would seem to be the only good way of verifying

that an actual experimental broadband position measurement reaches the "quantum regime" defined by equation (4.12)—unlike the scenario of optical measurements of quadrature amplitudes, there is no way directly to compare signal and meter beams to determine the fidelity of the measurement. If there is external force, the subtraction will not be perfect and α will reveal itself in nonstationary behavior of η_r.

Note as well that the variance of the residual noise process $\sqrt{\frac{C\sigma}{2}}\mathcal{N}_r[0,1]$ can in principle go to *zero* in the limit of good measurements. At first thought this might seem to have profound implications for the detection of weak classical forces, but in fact the quantitative analysis below recovers the usual SQL for weak force detection.

Before proceeding any farther, let us make the transition to the continuous limit ($\tau \to 0$, $\sigma \to \infty$, $D = \text{const}$) [71] and collect together a complete set of SDE's for the physical quantities and statistical estimators:

$$dx = \frac{p}{m}\,dt, \qquad dp = \alpha dt + \sqrt{\frac{\hbar^2}{2D}}dW, \tag{4.19}$$

$$d\xi = x dt + \sqrt{\frac{D}{2}}dW, \tag{4.20}$$

$$d\hat{x} = \frac{\hat{p}}{m}dt, \qquad d\hat{p} = \frac{\hbar}{D}\left(d\xi - \hat{x}dt\right). \tag{4.21}$$

In these terms, $d\xi/dt$ represents the broadband measurement record. Note that we could also define an auxilliary variable $e \equiv x - \hat{x}$, in terms of which

$$de = \frac{1}{m}\left(p - \hat{p}\right)dt, \tag{4.22}$$

$$d\hat{p} = \frac{\hbar}{D}\,e\,dt + \sqrt{\frac{\hbar^2}{2D}}dW. \tag{4.23}$$

Our estimator for the external force α will be proportional to the time integral of $d\eta \equiv d\xi - \hat{x}dt$, so what we really need to know is the rms behavior of $e\,(t)$. Developing its SDE,

$$\begin{aligned} de &= dx - d\hat{x}, \\ \frac{de}{dt} &= \frac{1}{m}\left(\alpha t - \frac{\hbar}{D}\int_0^t e\,(t')\;dt'\right), \end{aligned}$$

$$\frac{d^2e}{dt^2} = \frac{1}{m}\left(\alpha - \frac{\hbar}{D}e\right). \tag{4.24}$$

This shows that $e(t)$ behaves like a harmonic oscillator, with mass m and natural frequency $\omega_0 = \sqrt{\frac{\hbar}{mD}}$, subjected to a constant external force α. So with the initial condtion $\hat{x} = x$, $\hat{p} = p$, we expect

$$e(t) = \frac{\alpha D}{\hbar}\left[1 - \cos\left(\sqrt{\frac{\hbar}{mD}}t\right)\right]. \tag{4.25}$$

Getting back to the time integral of $d\eta$,

$$\int_0^t \left(\frac{d\xi}{dt'} - \hat{x}\right)dt' = \int_0^t e(t')\,dt' + \sqrt{\frac{D}{2}}\int_0^t dW(t') \tag{4.26}$$

from which we identify the signal Σ and rms noise N as

$$\Sigma = \frac{\alpha D}{\hbar}\left[t - \sqrt{\frac{mD}{\hbar}}\sin\left(\sqrt{\frac{\hbar}{mD}}t\right)\right], \quad N = \sqrt{\frac{Dt}{2}}. \tag{4.27}$$

Setting $\Sigma/N = 1$ we obtain

$$\alpha_{\min} = \hbar\sqrt{\frac{t}{2D}}\left[t - \sqrt{\frac{mD}{\hbar}}\sin\left(\sqrt{\frac{\hbar}{mD}}t\right)\right]^{-1}. \tag{4.28}$$

This expression gives a rigorous lower bound for the weakest detectable force, given particular values of m, t, and D. In order to arrive at an SQL in the traditional sense, what remains is to optimize over D. For fixed m and t, it seems reasonable that the optimal choice of D should be simply related to $\hbar t^2/m$. If we substitute $D \equiv \theta^2\hbar t^2/m$ into (4.28) and then (numerically) minimize over θ,

$$\alpha_{\min} = \min_{\theta}\sqrt{\frac{\hbar m}{2t^3}}\left[\theta - \theta^2\sin(1/\theta)\right]^{-1} \approx \pi\sqrt{\frac{\hbar m}{2t^3}}. \tag{4.29}$$

Hence we are able to recover (within a factor of π) the usual SQL for detecting a

weak classical force via repeated position measurements on a free mass [5].

Looking back at equations (4.14-4.15), note that we are allowed to substitute an *arbitrarily time-varying* external force $\alpha(t)$ without violating the conditions of the derivation (gaussian propagation). Hence, one arrives at the obvious extension of equation (4.24):

$$\frac{d^2e}{dt^2} = \frac{1}{m}\left(\alpha(t) - \frac{\hbar}{D}e\right). \tag{4.30}$$

The time-development of the Kalman estimator $e(t)$ therefore corresponds to the response of a simple harmonic oscillator (SHO) to the applied force $\alpha(t)$. All the quantum measurement phenomenology seems to be distilled into the single fact that the stiffness (resonance frequency) of this SHO scales inversely with the inaccuracy of the broadband position measurements. Computing this response gives you the exact waveform $e(t)$ that would be induced by any expected signal. Knowing this, one could easily design a matched filter [75] for optimal detection against the gaussian white noise $\sqrt{Dt/2}$.

Considering the simplest case $\alpha(t) = \alpha\cos(ft)$, we could guess that an interesting response might be obtained by choosing D such that $f = \sqrt{\hbar/mD}$. Under these conditions we can guess that $e(t) = A(t)\sin(ft)$, so that

$$\begin{aligned} \frac{d^2e}{dt^2} &= \left(\frac{\alpha}{m}\cos(ft) - f^2A(t)\sin(ft)\right) && (4.31)\\ &= 2f\dot{A}\cos(ft) + \left(\ddot{A} - f^2A\right)\sin(ft). && (4.32) \end{aligned}$$

Therefore, $A(t) = \alpha t\,(2mf)^{-1}$, and

$$\begin{aligned} e(t) &= \frac{\alpha t}{2mf}\sin(ft), && (4.33)\\ f &= \sqrt{\frac{\hbar}{mD}}. && (4.34) \end{aligned}$$

The matched filter for this will be simply $\sin(ft)$, and

$$\int_0^t dt' e(t')\sin(ft') = \frac{\alpha}{2mf^3}\int_0^t d(ft')\,(ft')\sin^2(ft')$$

$$= \frac{\alpha}{2mf^3}\left[\frac{1}{4}\sin^2(ft) - \right. \tag{4.35}$$

$$\left. \frac{ft}{2}\sin(ft)\cos(ft) + \frac{1}{4}f^2t^2\right] \tag{4.36}$$

The interesting regime would seem to be $ft \gg 1$, where the integrated error signal goes like t^2 (as opposed to t in the case of constant α). Setting Σ to the expression above, and $N = \frac{1}{2}\sqrt{Dt/2}$, we have

$$\alpha_{\min} \sim 4mft^{-2}\sqrt{Dt/2} \tag{4.37}$$

$$= 4\sqrt{\frac{\hbar m}{2t^3}}. \tag{4.38}$$

Although the derivation of this last expression has not been completely rigorous, one sees that the use of equation (4.30) together with some intuition about the behavior of forced harmonic oscillators provides a convenient way of thinking about optimal signal recovery in the quantum regime of strong measurement backaction.

4.2 Quantum system identification

Abstract: The rapidly-developing theory of quantum parameter estimation (QPE) focuses on the design of optimal measurement strategies for extracting information about c-number parameters $\vec{\theta}$ that characterize a given quantum system. While much progress has been made in applying QPE to the parametric identification of *prepared quantum states*, little or no attention has yet been paid to the problem of estimating parameters that characterize a *dynamical quantum system.* In this section I consider the latter aspect of QPE within the context of quantum optics, and describe a quantum trajectories method for estimating parameters that appear in the effective Hamiltonian for a Markovian open quantum system.

Recent theoretical work in QPE [76, 77, 15] has focused on a paradigm in which an experimenter is provided with one or more copies of a quantum state ρ_0 drawn from a single-parameter family $\rho(\theta)$, and is asked to determine the value θ_0 such

that $\rho_0 = \rho(\theta_0)$. The experimenter knows the form of $\rho(\theta)$ and can make arbitrary measurements on the states she is given, but does not know the value of θ_0 *a priori*. In such situations one can actually derive a mathematical representation of the optimal quantum measurement for the purpose of estimating θ_0, and optimize over all possible statistical reductions of the measurement results as well [78]. Accordingly, there exists a generalized version of the Cramér-Rao inequality [79] that establishes a fundamental bound on the rate of convergence for estimators based on repeated measurements whose marginal statistics are fully determined by a single, unchanging density matrix [80]. My purpose here is to consider a related but distinct aspect of QPE, namely the estimation of parameters appearing in the equations of motion that govern the time-evolution of a quantum system. In this paradigm, which is closely related to that of classical *system identification* [81], the hypothetical experimenter wishes to determine which system model $\mathcal{H}_0$ in a parametrized family $\mathcal{H}(\vec{\theta})$ best accounts for the dynamical behavior of a given quantum system. In contrast to the conventional QPE scenario described above, the statistics of multiple successive measurements made on a dynamical system cannot necessarily be derived from any single density matrix. Roughly speaking, this is because every measurement disturbs the state of the system [12] in a manner that depends on the measurement outcome, and because the evolution of the system state between measurements depends on the full details of the system's equations of motion.

The effects of repeated measurements on otherwise-unitary quantum evolution have been extensively studied in quantum optics, with regard to the dynamics of open quantum systems [9, 8]. The configuration most often treated by such work is that of a small, "encapsulated" quantum system having one or more well-defined input/output channels associated with its coupling to the physical environment. This picture naturally suggests a paradigm in which an experimenter attempts to parametrically identify the Hamiltonian of the encapsulated quantum system by examining the response of the output channels to driving stimuli applied to the input channels [82]. *The task of quantum system identification may then be equated with that of computing the relative likelihood of an observed sequence of measurement results* Ξ^*

as a function of the parameter set $\vec{\theta}$, given the external driving conditions imposed by the experimenter. To the extent that the environmental couplings for the system are known, quantum trajectory theory [8, 13] suggests a simple method for the computation, which I discuss below. Having a likelihood function $f(\vec{\theta}\,|\,\Xi^*)$, one can use maximum-likelihood or Bayesian principles [83, 81] to estimate the parameters $\vec{\theta}$. Note that it should generally be possible for the experimenter to determine optimal driving conditions that make the system response maximally sensitive to the values of $\vec{\theta}$, or indeed to *adaptively* change the driving conditions as the estimation starts to converge [15].

To illustrate quantum system identification in a concrete setting, let us focus on an example with relevance to current experiments in cavity quantum electrodynamics (QED)—a single two-level atom placed within the mode volume of a driven, high-finesse optical cavity [22, 8]. The strength of the coherent coupling between atom and cavity mode is parametrized by the vacuum Rabi frequency g, whose value depends on the spatial position of the atom within the cavity. For a Fabry-Perot resonator $g(\vec{r}) = g_0 \cos(2\pi x/\lambda) \exp\left[-(y^2+z^2)/w_0^2\right]$, where x is the coordinate along the cavity axis and w_0 is the gaussian waist of the TEM_{00} resonator mode. The specific task I shall consider is that of estimating $g \in [0, g_0]$, which I suppose to be unknown because the atomic position is not known. The measurement procedure will simply be to monitor the arrival-time statistics of photons emitted by the atom-cavity system for a fixed cavity driving field. For the purposes of this discussion I shall not explicitly treat the atom's external degrees of freedom, imagining that they are fixed by an rf Paul trap or similar confining mechanism [84]. However, note that the correlation of g with the atomic position *operator* implies that "online" estimation of g for an untrapped atom drifting through a cavity could be viewed as a time-distributed quantum measurement of the position of a free mass [85, 86].

For a *gedanken*-experiment in which the cavity is driven by a resonant cw probe laser and both the atomic fluoresence and cavity emission are continuously monitored by perfect photon-counting detectors [87], the evolution of the conditional state-vector

between photodection events satisfies the effective Schrödinger equation ($\hbar = 1$)

$$|\psi_c(t+dt)\rangle = e^{-i\mathcal{H}_{\rm eff}dt}|\psi_c(t)\rangle, \tag{4.39}$$

$$\mathcal{H}_{\rm eff}(g) = ig\left(a\sigma_+ - a^\dagger\sigma_-\right) + i\epsilon\left(a - a^\dagger\right) - i\kappa a^\dagger a - i\gamma_\perp\sigma_+\sigma_-. \tag{4.40}$$

This interaction-picture expression for $\mathcal{H}_{\rm eff}(g)$ is valid under the rotating-wave and electric-dipole approximations, and for identical atomic/cavity/probe-laser frequencies. Here κ is the field decay rate of the cavity, $\gamma_\perp$ is the dipole decay rate of the atom, and ϵ represents the strength of the coherent driving field. The jump operator associated with the detection of photons spontaneously emitted by the atom is $\hat{c}_0 \equiv \sqrt{2\gamma_\perp}\sigma_-$, and $\hat{c}_1 \equiv \sqrt{2\kappa}a$ is the jump operator associated with the detection of photons leaking through the cavity mirrors.

By registering the origins $\{j_1, \ldots, j_n\}$ (= 0 for spontaneous emission or 1 for cavity decays) and arrival times $\{t_1, \ldots, t_n\}$ of every photon emitted by the atom-cavity system in response to the cw driving field during an observation interval $[t_0, t_f]$, the hypothetical experimenter accumulates a classical record $\Xi^* \equiv (t_0, t_f, \{j_i, t_i\})$ of the stochastic evolution of the system state. Assuming a uniform prior distribution on g, the likelihood function $f(g \,|\, \Xi^*)$ then simply corresponds to a normalized version of the exclusive probability density [8, 13]

$$p\left(\Xi^* | \, g\right) = \mathrm{Tr}\left[\begin{array}{c} U_{\rm eff}\left(t_f, t_n\,;g\right)\hat{c}_{j_n}U_{\rm eff}\left(t_n, t_{n-1}\,;g\right)\cdots \\ \times\,\hat{c}_{j_1}U_{\rm eff}\left(t_1, t_0\,;g\right)\,\rho(t_0)\,\,U^\dagger_{\rm eff}\left(t_1, t_0\,;g\right)\,\hat{c}^\dagger_{j_1} \\ \times\,\cdots U^\dagger_{\rm eff}\left(t_n, t_{n-1};g\right)\,\hat{c}^\dagger_{j_n}\,U^\dagger_{\rm eff}\left(t_f, t_n;g\right) \end{array}\right], \tag{4.41}$$

viewed as a function of g rather than Ξ^*. Accordingly, the maximum-likelihood estimate (MLE) of g is obtained by computing the value of g which maximizes (4.41) with Ξ^* fixed by the observed data. Here $U_{\rm eff}\left(t', t\,;g\right)$ is the evolution operator from time t to t' associated with the effective Hamiltonian $\mathcal{H}_{\rm eff}(g)$ defined in equation (4.40).

In order to numerically demonstrate quantum system identification using (4.41), I have generated a set of classical records by quantum Monte Carlo simluation [88] of a driven atom-cavity system with $(g_0, \gamma_\perp, \kappa)/2\pi = (57, 2.5, 30)$ MHz, and three

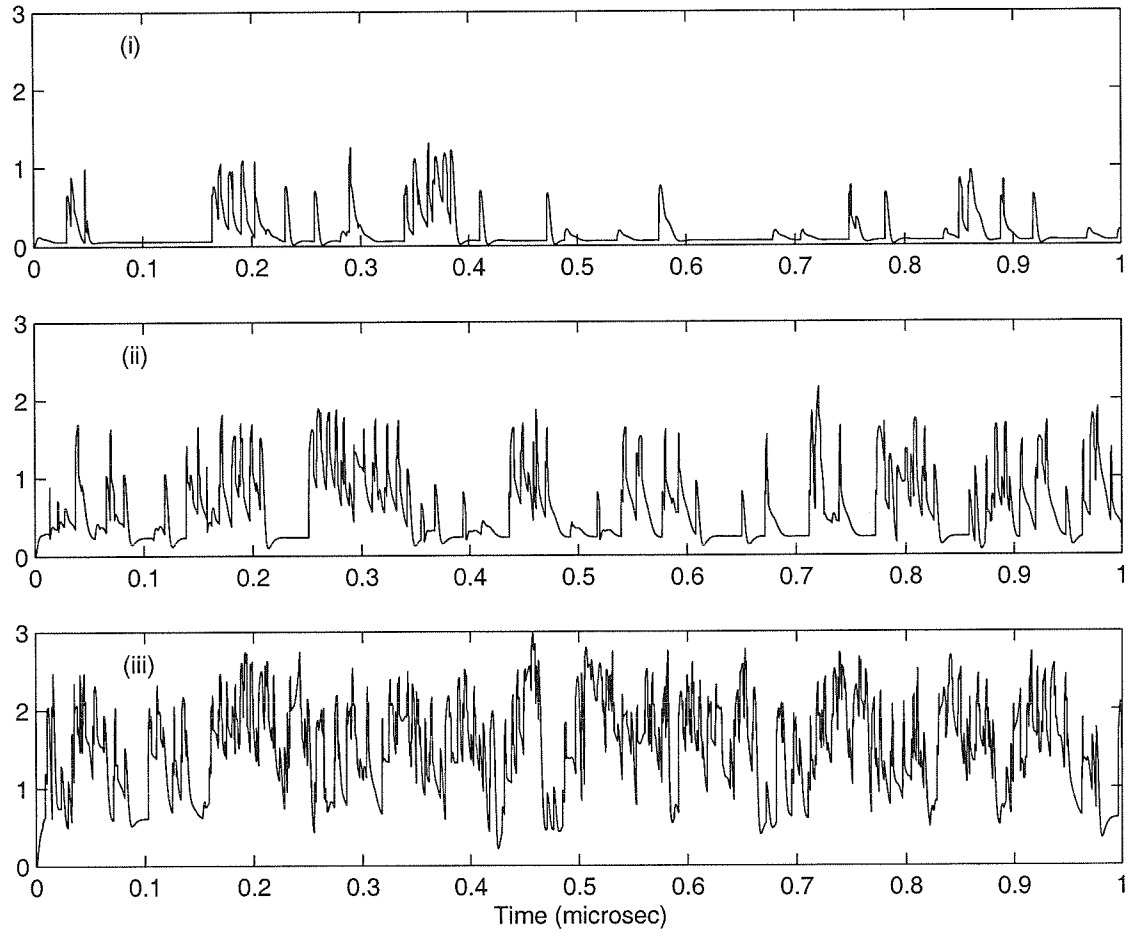

Figure 4.2: Time-evolution of the mean intracavity photon number $\langle a^\dagger a \rangle$ in individual trajectories. Top trace (i) is for $\epsilon = 24$, middle trace (ii) is for $\epsilon = 34$, and bottom trace (iii) is for $\epsilon = 44.3$.

different powers for the driving field $\epsilon = \{24, 34, 44.3\}$. While this value for g_0 is larger than what has been achieved experimentally [89], it should certainly be within reach of works in progress. For the simulations I chose an arbitrary atomic position such that $g(\vec{r}) = 45$ MHz, and generated classical records with an observation time of 1 μs each. Figure 4.2 illustrates the stochastic time-evolution of the mean intracavity photon number, taken from typical Monte-Carlo data sets for each of the three values of ϵ. The photocount statistics are clearly super-Poissonian, and the simulated data show that quantum jumps often occur at a local rate that greatly exceeds the rate at which the system regresses to steady state.

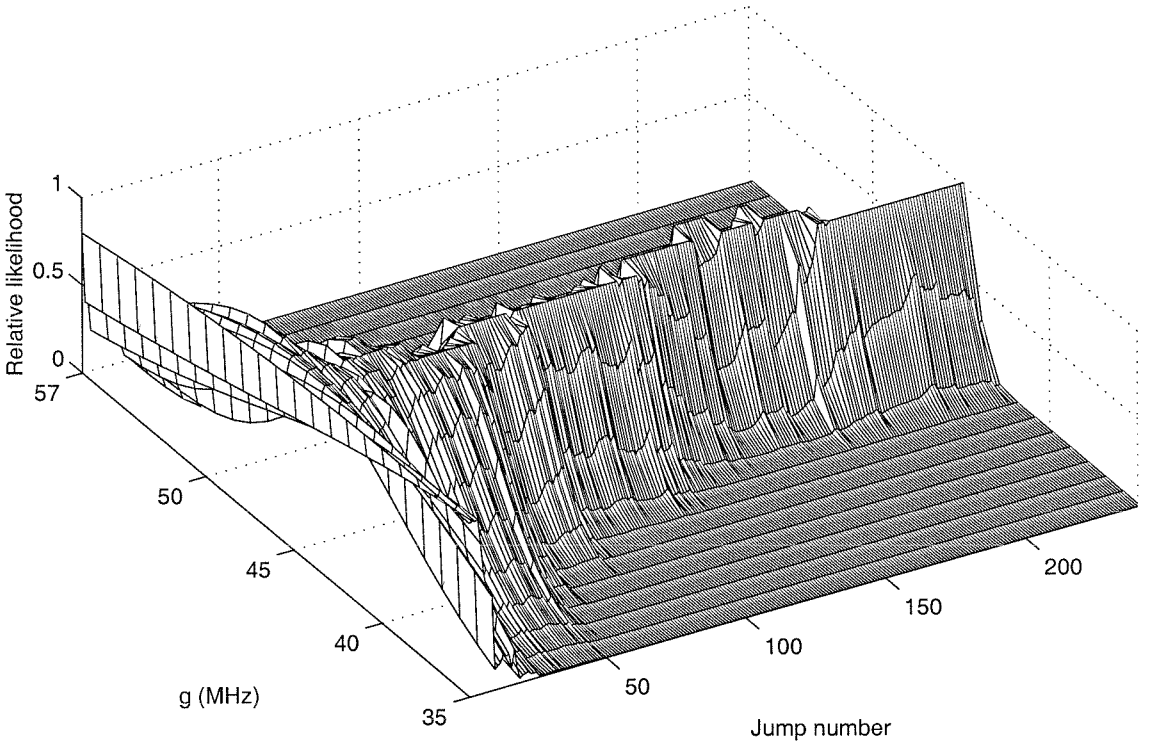

Figure 4.3: Corresponding stochastic evolution of the (normalized) likelihood function $f(g|\Xi^*)$ and corresponding MLE in one quantum trajectory with driving field amplitude $\epsilon = 34$. The surface height indicates relative probability of $g \in [35, 57]$, with a resolution of 1. The "true" value of g corresponds to 45. Note that the likelihood function is updated each time a photon is detected, so that the timelike coordinate in this surface plot corresponds to jump number rather than absolute time.

For each Monte Carlo trajectory, an identification routine based on (4.41) was used to compute the stochastic time-evolution of $f(g\,|\,\Xi^*)$, as well as the corresponding MLE. Figure 4.3 shows one typical data set with $\epsilon = 34$, starting from the inital estimate made after only one photodetection event and updated after each subsequent photodetection. Figure 4.4 indicates the ensemble-averaged convergence of the MLE for g, based on 2000 simulations for $\epsilon = 24$, 300 simulations for $\epsilon = 34$, and 150 simulations for $\epsilon = 44.3$. A histogram representing the time-evolution of the MLE sampling distribution is given in Figure 4.5, for the case of $\epsilon = 44.3$. With this driving field, $\sim 1\%$ accuracy in estimation of g is obtained in 1 μs observation time (~ 600 quantum jumps).

It is important to note that the QPE procedure described above automatically

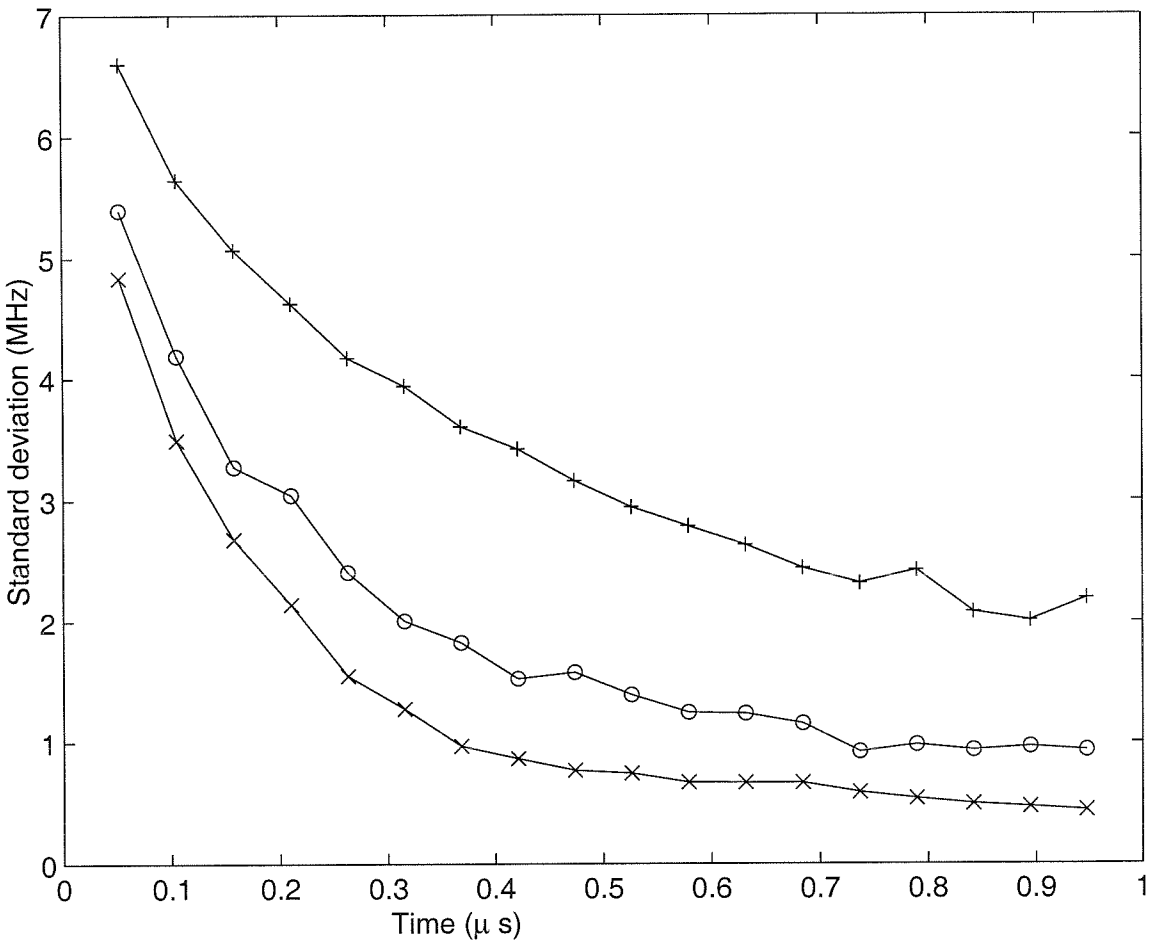

Figure 4.4: Standard deviation of the maximum-likelihood estimator for g as a function of absolute time (+—$\epsilon = 24$, ○—$\epsilon = 34$, ×—$\epsilon = 44.3$).

makes efficient use of any information about g that is contained in higher-order correlations of the classical record of counting times. Of course, not every open quantum system will generate significant correlations of this type. In the scenario discussed above for example, the photon stream emitted by the atom-cavity system would become nearly Poissonian in the limit of either weak excitation (corrleated pairs of photons become rare) or of weak coupling $g \ll \kappa, \gamma_\perp$ (correlations become weak). Two critical conditions for correlations to be strongly evident in individual classical records are that the mean time between counts must be comparable to or less than the system regression time [93], and that the system dynamics must be significantly altered by the loss of a single quantum of excitation [90]. For systems not satisfying

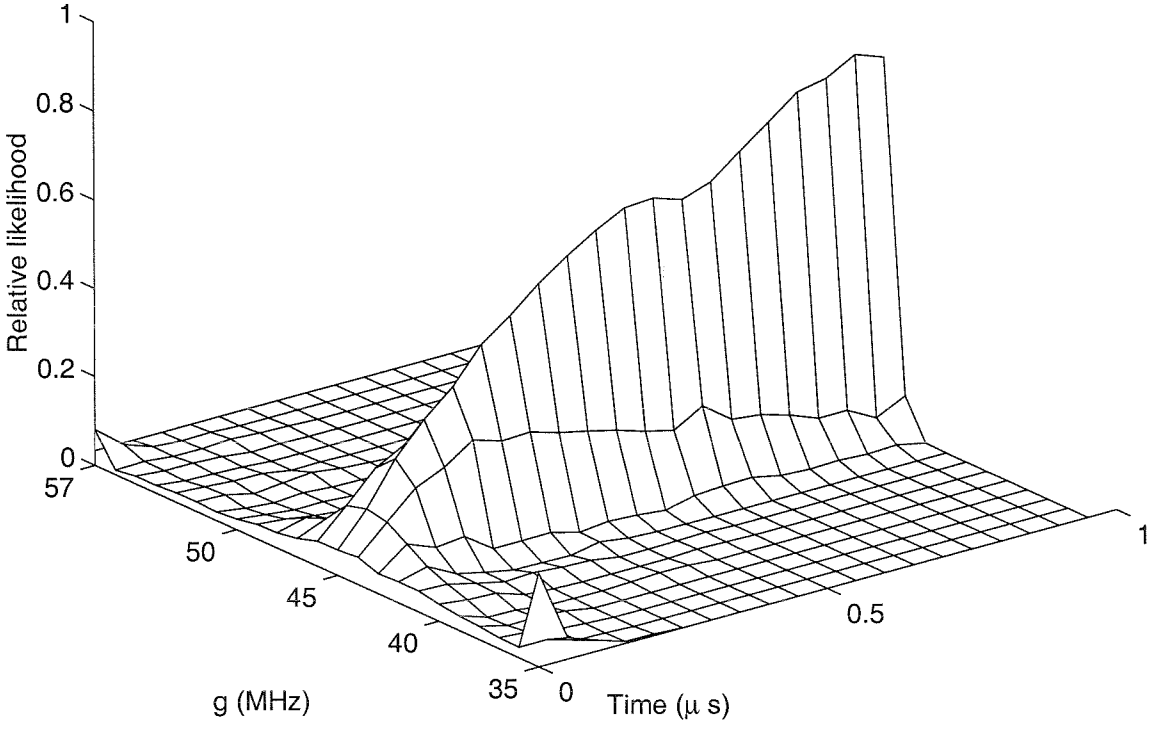

Figure 4.5: Histogram showing the evolution (in absolute time) of the sampling distribution for the MLE of g, representing 150 simulations with $\epsilon = 44.3$.

these criteria, the methods described above offer no real advantage over statistical estimators based on only the *steady-state* density matrix obtained by solving the master equation associated with (4.40). Accordingly, optimal parameter estimation in such systems can be formulated within the paradigm of conventional QPE. Judging from the trend shown in Figure 4.4, however, it certainly seems that the information rate on parameters in strongly-coupled systems can be significantly larger in the strong-driving regime [91, 92] than in the weak-field regime.

In closing, let me note that a straightforward extension of the above method would allow the identification of non-stationary Hamiltonians in which the parameters $\vec{\theta}(t)$ vary slowly compared to the timescale for convergence of the corresponding statistical estimators. Relative to the example discussed above, a recent cavity-QED experiment incorporating a laser-cooled atomic source [26] has demonstrated the practical fea-

sibility of achieving this separation of timescales. It seems reasonable to hope that the methods proposed above could be utilized in future experimental work to track variations in g associated with the motion of an individual atom through the mode-volume of a high-finesse optical cavity [87]. A digital signal processor implementing such a procedure could be used as the state observer in a "semiclassical" feedback control loop designed to confine and cool the atom's center-of-mass motion.

Appendix A Semiclassical measurement strategies

The notes in this Appendix describe some statistical methods for optimal estimation of intracavity atom number, based on the photocount statistics of the cavity output. All derivations are based on one *very* incorrect assumption, however, namely that the cavity output remains in a coherent state regardless of the value of the cooperativity C. After realizing that this was a fatal problem, I set off down the road that led to the quantum system identification paper. So although the work presented below turned out to be nothing more than an edifying exercise, I think it's worth preserving as a general illustration of how to apply classical parameter estimation methods in quantum optics.

A.1 Preliminaries

We consider the difference between atom *counting* and atom *detection* from a statistical point of view. In the case where one wants maximum resolution for determining the precise number of atoms in a cavity, the appropriate figure-of-merit in parameter space should be Fisher information on the cooperativity C. Hence the applicable statistical framework is that of point-estimation theory, and in particular one knows to search for an *efficient estimator* of C. If on the other hand one is interested only in distinguishing between the presence and absence of atoms in a cavity, hypothesis-testing theory should be applied. Then the goal is to find a *uniformly-most-powerful test* of the hypothesis that the cooperativity is greater than some threshold C_0. To illustrate the distinction we derive optimal parameters for photon-counting measurements of cavity transmission, with one atom in the bad-cavity limit and an on-resonance probe.

For a single atom interacting with a quantized cavity mode, Rice and Carmichael

[93] have derived an expression for the steady-state intracavity intensity

$$\langle a^\dagger a \rangle_{ss} = m_0 Y^2 \left(\frac{1+Y^2}{\left[1+2C(\vec{\mathbf{r}})\right]^2 + Y^2} \right), \tag{A.1}$$

which is valid in the bad-cavity limit for a resonant driving field of arbitrary power Y. Here $\vec{\mathbf{r}}$ denotes the atom's position inside the cavity. We restrict our attention to this case because it's nice to have an actual function with which to do analytic manipulations (as opposed to the unrestrained bistability state equation). With only one atom in the cavity the distinction between counting and detection is not so sharp, but we'll take atom counting to mean "precise determination of $C(\vec{\mathbf{r}})$," and atom detection to mean "deciding whether $C(\vec{\mathbf{r}}) \geq C_0$." Note that for our scenario $0 \leq C(\vec{\mathbf{r}}) \leq C_1$, where $C_1 \equiv g_0^2/2\kappa\gamma_\perp$ is the single-atom cooperativity. Writing $C(\vec{\mathbf{r}}) \equiv C_1 \Phi(\vec{\mathbf{r}})$, we could just as well think of our real goal as being estimation of the "position parameter" $\Phi(\vec{\mathbf{r}}) \in [0,1]$.

The photon-counting rate for the transmitted field depends upon $C \equiv C(\vec{\mathbf{r}})$ through the Poissonian mean

$$\Gamma = \frac{2\kappa m_0}{B} Y^2 \left(\frac{1+Y^2}{1+4C+4C^2+Y^2} \right), \tag{A.2}$$

where B is the detection bandwidth, 2κ the energy decay rate of the cavity, and m_0 the critical photon number. In the following we will need the first two derivatives of Γ with respect to C,

$$\frac{\partial \Gamma}{\partial C} = -\frac{8\kappa m_0 Y^2}{B} \frac{(1+Y^2)(1+2C)}{(1+4C+4C^2+Y^2)^2}, \tag{A.3}$$

$$\frac{\partial^2 \Gamma}{\partial C^2} = \frac{16\kappa m_0 Y^2 (1+Y^2)}{B} \left[\frac{3+12C+12C^2-Y^2}{(1+4C+4C^2+Y^2)^3} \right]. \tag{A.4}$$

An excellent reference for the statistical theory discussed below is [83].

A.2 Atom counting

In this section we consider the task of trying to determine the precise value of C through photon-counting measurements of the cavity output. The number of photocounts received in a given detection bandwidth is a Poissonian random variable whose mean is tied to the intracavity photon number, and therefore to C. Having an expression for the probability distribution for different numbers of photocounts, we can compute a quantity called the *Fisher Information* on C, denoted $I(C)$. This quantity holds great significance for our given task because the *Cramèr-Rao bound* states that no statistical estimator of C can converge with variance smaller than $1/nI(C)$, where n is the number of photocount "bins" we integrate over. The information $I(C)$ will turn out to be a function of the driving field Y, such that we can derive an *optimum* value of Y for determining C from the photon-counting distribution.

The likelihood (*i.e.*, probability distribution) to observe a given number of counts r in one detection bandwidth is Poissonian,

$$f(r\,|\,C) = \frac{e^{-\Gamma}\Gamma^r}{r!}\,. \tag{A.5}$$

Taking derivatives of the log-likelihood $\lambda(r\,|\,C)$, we find

$$\lambda'(r\,|\,C) = -\left(\frac{\partial\Gamma}{\partial C}\right) + \frac{r}{\Gamma}\left(\frac{\partial\Gamma}{\partial C}\right), \tag{A.6}$$

$$\lambda''(r\,|\,C) = -\left(\frac{\partial^2\Gamma}{\partial C^2}\right) - \frac{r}{\Gamma^2}\left(\frac{\partial\Gamma}{\partial C}\right)^2 + \frac{r}{\Gamma}\left(\frac{\partial^2\Gamma}{\partial C^2}\right). \tag{A.7}$$

Using the expressions given above for derivatives of Γ, we arrive at the Fisher information on C,

$$I(C) = -E\left[\lambda''(r\,|\,C)\right] = \frac{1}{\Gamma}\left(\frac{\partial\Gamma}{\partial C}\right)^2, \tag{A.8}$$

$$= \frac{32\kappa m_0}{B}\frac{Y^2\left(1+Y^2\right)\left(1+2C\right)^2}{\left(1+4C+4C^2+Y^2\right)^3}\,. \tag{A.9}$$

Having this expression, we can maximize (at least numerically) with respect to Y

and determine the optimal driving field. Note that the optimal Y will certainly vary with C, so that one can only really maximize the information for a given expected value. The goal then is to find an *efficient estimator* for C. There isn't really a general method for deriving such an estimator; in fact, there doesn't even seem to be a test for the existence of an efficient estimator. However, as the sample mean $\bar{r}$ is a *sufficient statistic* for C in the Poisson distribution (A.5), it seems reasonable just to invert the expression for $E[\,r\,|\,C\,]$. This procedure yields the "phenomenological" estimator

$$\hat{C} = \frac{1}{2}\left[-1 \pm \sqrt{\frac{2\kappa m_0 Y^4}{\bar{r}B} + \left(\frac{2\kappa m_0}{\bar{r}B} - 1\right) Y^2}\right] . \tag{A.10}$$

Thus defined, $\hat{C}$ is an unbiased estimator of C. To show that it is in fact efficient, we would need to either compute its variance, or else show that it can be factored in the form

$$\hat{C} = u(C)\,\lambda_n'(\,\{r_i\}\,|\,C\,) + v(C)\,, \tag{A.11}$$

where

$$\begin{aligned} \lambda_n'(\,\{r_i\}\,|\,C\,) &= (-1 + n\bar{r})\,\frac{\partial \Gamma}{\partial C}\,, \\ &= (1 - n\bar{r})\left(\frac{8\kappa m_0 Y^2}{B}\,\frac{(1+Y^2)\,(1+2C)}{(1+4C+4C^2+Y^2)^2}\right). \end{aligned}$$

Unfortunately, the algebra gets rather cumbersome so we'll have to leave this issue unsettled for now. At any rate, the Cramèr-Rao bound tells you that no estimator of C can converge with variance smaller than $1/nI(C)$, where n is the number of detection times you integrate over. In fact it can be shown that the *Maximum Likelihood Estimator* of C should saturate the Cramèr-Rao bound *asymptotically* (that is, for large n). It would therefore be interesting to derive its form (coming soon).

In order to get a qualitative feel for what's going on, we show in Figure A.1 the Fisher information on C as a function of driving field Y. These plots show the optimum driving field for precise determination of an unknown but fixed cooperativity C. Note that a slightly stronger maximum exists for $C = 1$ than for higher values.

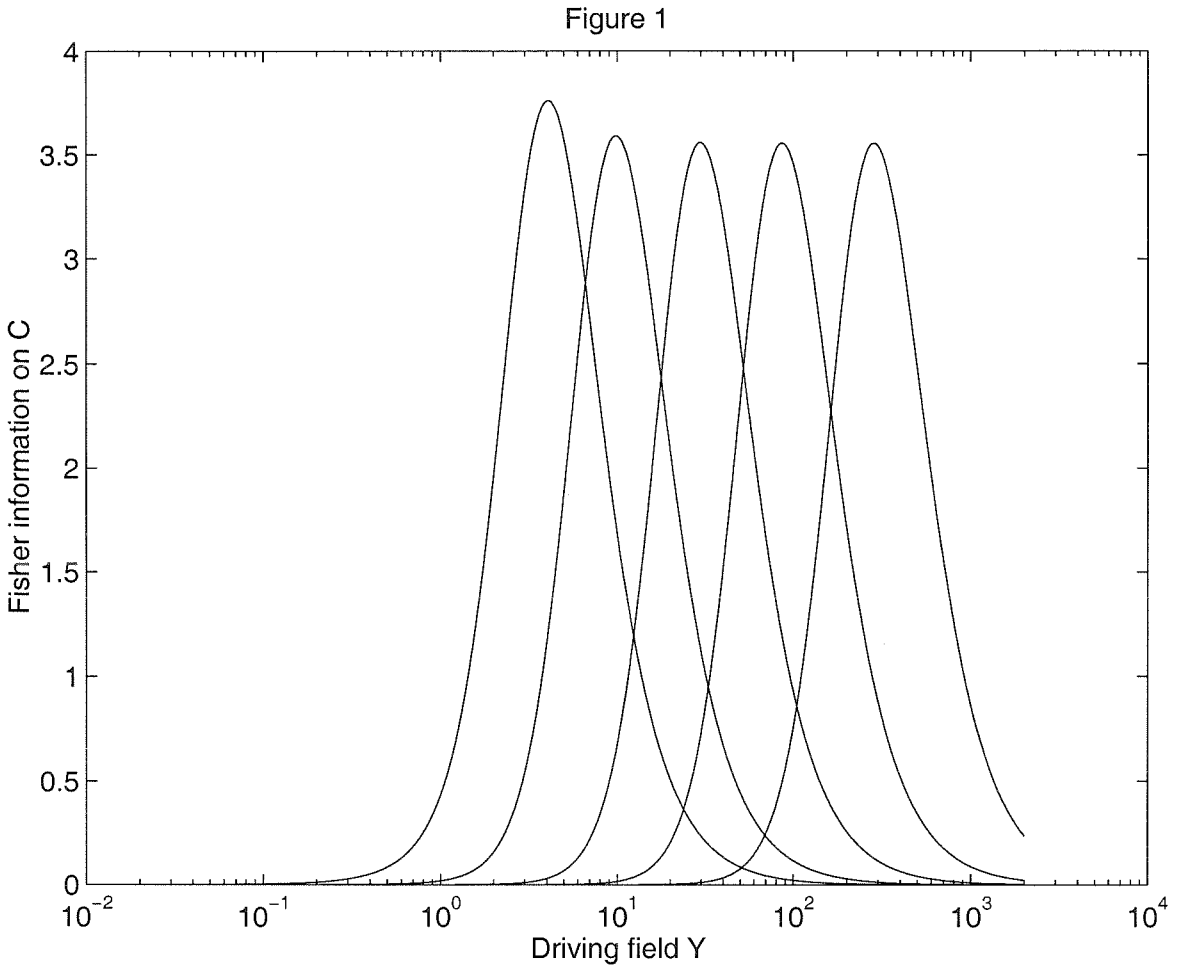

Figure A.1: Fisher information on C versus Y, for C=1, 3, 10, 30, and 100 (appearing left-to-right in the plot).

Figure A.2 shows $I(C)$ for fixed driving fields Y and varying "actual" cooperativity. The plots indicate that with too weak a driving field, there's hardly any information about C. For driving fields at or above saturation, good information is obtained for some range of cooperativities whose center value increases with Y. Keep in mind that the axes are logarithmic, so that larger driving fields also provide information for a wider range of cooperativities.

A.3 Atom detection

We next turn to the more discrete task of deciding whether C is less than or greater than some "detection threshold" C_0. For this task we must set aside the Cramèr-Rao

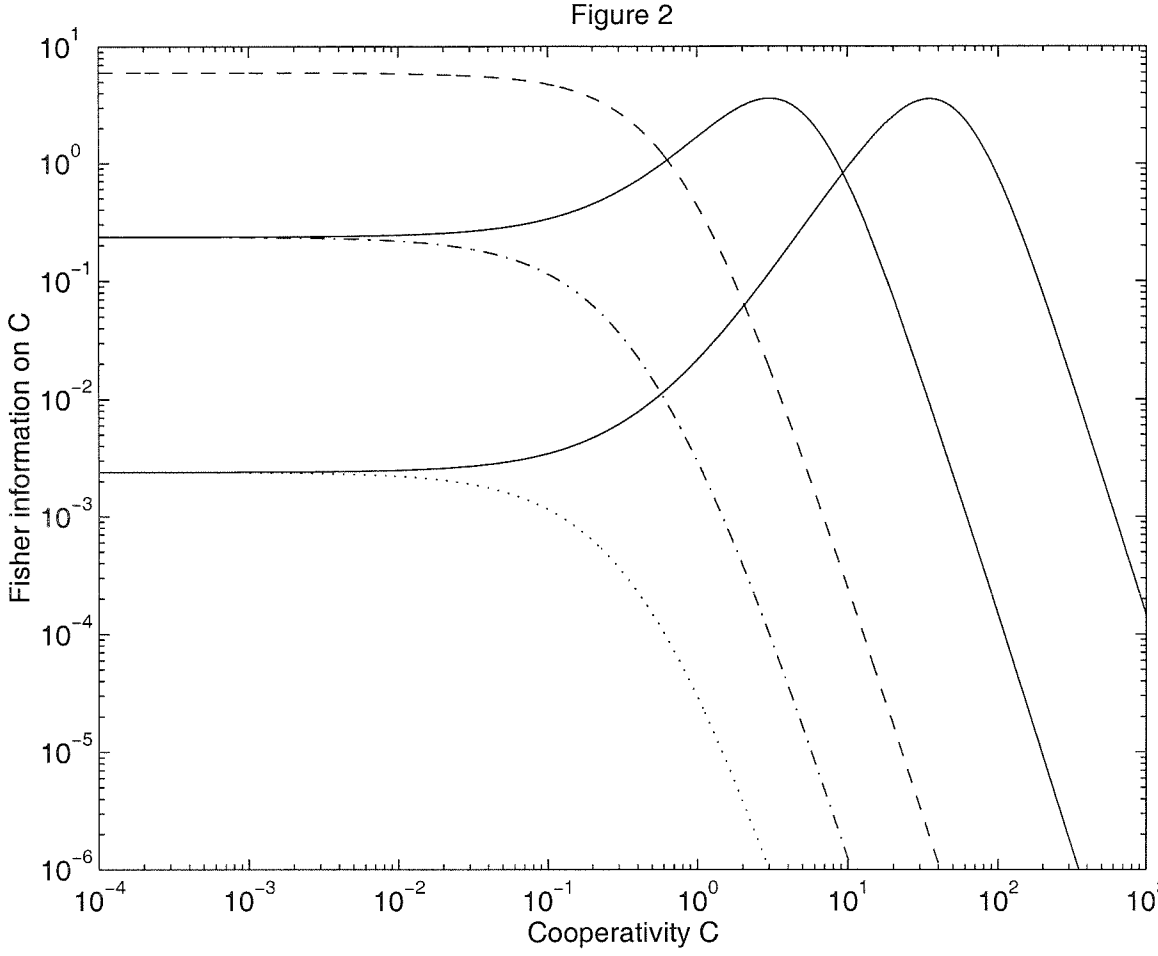

Figure A.2: Fisher information on C versus C, for Y=0.01 (dotted), 0.1 (dot-dash), 1 (dashed), 10 (solid), and 100 (solid).

bound in favor of something called the *Neyman-Pearson lemma* for hypothesis-testing. We can cast our atom detection problem in the paradigm of *one-sided alternatives* by formulating two hypotheses about C:

$$H_o \;:\; C \leq C_0 \,, \tag{A.12}$$

$$H_1 \;:\; C > C_0 \,. \tag{A.13}$$

Here H_0 is the "null hypothesis" that no atoms are strongly-coupled to the cavity (*i.e.*, $C(\vec{\mathbf{r}}) < C_0$), and H_1 represents rejection of the null hypothesis. By fixing an upper limit α_0 on the false-alarm probability (*i.e.*, a level of significance), we can use the following condition (a consequence of Neyman-Pearson) to derive an optimal

testing procedure (a *uniformly-most-powerful test*):

$$\Pr(\, T \geq c \,|\, C = C_0 \,) = \alpha_0. \tag{A.14}$$

Here T represents a special statistic for which $f(\, r \,|\, C \,)$ has a monotone likelihood ratio. Knowing the sampling distribution of T, we are meant to choose the constant c so as to satisfy this condition. Then the optimal testing procedure consists of rejecting H_0 if and only if $T \geq c$. The figure-of-merit with which we should optimze Y is the integral (or average) of the power function $\pi(C)$ (defined as the probability of rejecting H_0, given that the actual value of the cooperativity is C) over the region $C_0 < C \leq C_1$. The value of this integral corresponds to the probability that a "real event" is successfully detected (given uniform expected distribution for C, a condition which could be relaxed), and should of course be maximized.

The computational difficulty lies in identifying a statistic T for which $f(\, r \,|\, C \,)$ has a monotone likelihood ratio. Specifically, we need the likelihood ratio

$$\frac{f_n(\, \{r_i\} \,|\, C_a \,)}{f_n(\, \{r_i\} \,|\, C_b \,)} = e^{-n(\Gamma_a - \Gamma_b)} \left(\frac{\Gamma_a}{\Gamma_b} \right)^{n\bar{r}} \tag{A.15}$$

to depend on the n data points $\{r_i\}$ only through a statistic T, and the value of this ratio must be a monotonically-increasing function of T for $C_a > C_b$. As we know that Γ is a strictly *decreasing* function of C, we know that $(\Gamma_a/\Gamma_b)^n < 1$ so there is actually an easy solution of just choosing $T \equiv (n\bar{r})^{-1}$. Using this we can then compute the constant c defined above, and then try to optimize the figure-of-merit

$$\Pi \equiv \frac{1}{C_1 - C_0} \int_{C_0}^{C_1} \pi(C)\, dC \tag{A.16}$$

$$= \frac{1}{C_1 - C_0} \int_{C_0}^{C_1} \Pr(T \geq c)\, dC \tag{A.17}$$

with respect to Y.

Carrying out this procedure, we first derive the sampling distribution for T. Con-

sider first the distribution of the "sample norm" $S \equiv n\bar{r}$,

$$f(s\,|\,n,C) = \sum_{n\bar{r}=s} f_n(\{r_i\}\,|\,C) \tag{A.18}$$

$$= \sum_{n\bar{r}=s} \frac{e^{-n\Gamma}\Gamma^{n\bar{r}}}{\prod_{i=1}^{n}(r_i!)} \tag{A.19}$$

$$= e^{-n\Gamma}\Gamma^{s} \sum_{n\bar{r}=s} \frac{1}{\prod_{i=1}^{n}(r_i!)}, \tag{A.20}$$

where the sum is over all sets of values $\{r_i\}$ such that $\sum_{i=1}^{n} r_i = s$. Transforming then to the variable $T = S^{-1}$, we have

$$f(T = t\,|\,n,C) = f(S = t^{-1}\,|\,n,C) \tag{A.21}$$

$$= e^{-n\Gamma}\Gamma^{1/t} \sum_{n\bar{r}=1/t} \frac{1}{\prod_{i=1}^{n}(r_i!)}\,. \tag{A.22}$$

Note that the "allowed" values of T are $t = \infty, 1, \frac{1}{2}, \frac{1}{3}, \ldots$ where inspection of expression (A.22) confirms that there is nothing pathological about the infinite value. At this point the only remaining obstacle to determining the constant c in (A.14) is computation of the entropy-like sum in (A.22). Unfortunately, there doesn't seem to be a nice analytic simplification, but for reasonable values of n it can certainly be tabulated.

The situation is simplified greatly if we restrict our attention to the case where $n = 1$. Indeed, the fact that the sample mean is a sufficient statistic for C suggests that we might not loose anything by increasing the integration time B^{-1} and testing the hypothesis (A.12) one data point at a time. Then the distribution (A.22) reduces to

$$f(T = t\,|\,C) = \frac{e^{-\Gamma}\Gamma^{1/t}}{(t^{-1})!}\,. \tag{A.23}$$

Inserting this into the test condition (A.14), we can finally derive an expression for the threshold constant c:

$$\sum_{t=c}^{\infty} \frac{e^{-\Gamma_0}\Gamma_0^{1/t}}{(t^{-1})!} = \alpha_0\,, \tag{A.24}$$

where

$$\Gamma_0 \equiv \frac{2\kappa m_0}{B} Y^2 \left(\frac{1+Y^2}{1+4C_0+4C_0^2+Y^2} \right) . \tag{A.25}$$

Having written this down, we can transform back to the less confusing variable $s \equiv t^{-1}$, which ranges over the natural numbers 0, 1, 2, ... Let us also define $c' \equiv c^{-1}$, so that (A.24) becomes

$$\sum_{s=0}^{c'} \frac{\Gamma_0^s}{s\,!} = \alpha_0 \, e^{\Gamma_0} . \tag{A.26}$$

Having a threshold constant c' defined in this way, the uniformly-most-powerful test of hypothesis (A.12) will be to decide that there is an atom in the cavity if and only if the number of counts s recieved in a detection bandwidth B is *less* than or equal to c'.

For fixed significance α_0 and threshold-cooperativity C_0, the test-condition constant c' is clearly still a function of Y. In order to determine the optimal driving field, we will also need to maximize the overall detection probability Π for $C_0 < C \leq C_1$. Recalling expression (A.17) from above,

$$\begin{aligned} \Pi(Y) &= \frac{1}{C_1 - C_0} \int_{C_0}^{C_1} \Pr(T \geq c)\, dC \,, \\ &= \frac{1}{C_1 - C_0} \int_{C_0}^{C_1} \Pr(S \leq c')\, dC \,, \\ &= \frac{1}{C_1 - C_0} \int_{C_0}^{C_1} dC \; e^{-\Gamma(C,Y)} \left[\sum_{s=0}^{c'} \frac{\Gamma^s(C,Y)}{s\,!} \right] , \end{aligned} \tag{A.27}$$

where the notation $\Gamma(C,Y)$ is meant to indicate the dependence of the expected count-rate on C and Y. After integrating out C in expression (A.27), the overall detection probability Π is a function only of Y. Hence by maximizing (A.27) with respect to Y we can determine the optimum driving field strength—optimum in the sense that the false-alarm probability is $\leq \alpha_0$ (by construction), and the probability of successfully detecting a real event is as large as possible. Note that this latter property was obtained in two steps: for *fixed* Y, the test-condition constant c' was determined according to the critereon (A.14), which guarantees that the test $T \geq c'$

has the *maximum* probability to succeed, *among all testing procedures of the form* $T \geq x$, for *every* value of $C > C_0$. Then the *variation* with respect to Y ensures that the *overall* detection probability, *integrated* over the possible values $C > C_0$, is also maximized.

Figure A.3 shows $\Pi(Y)$ for an example scenario, with $C_1 = 22$, $C_0 = \{1, 5, 10\}$, and $\alpha = 10^{-5}$, for which the test-condition constants turn out to be 54, 126, and 129, respectively. In Figure A.4 we show the power function $\pi(C)$ for the particular case of $C_0 = 5$. For comparison we show in Figure A.5 the Fisher information on C at the optimum values of Y determined by the hypothesis-testing criterion. And for completeness, Figure A.6 shows the Fisher information on C as a function of Y, for $C = \{1, 5, 10, 22\}$.

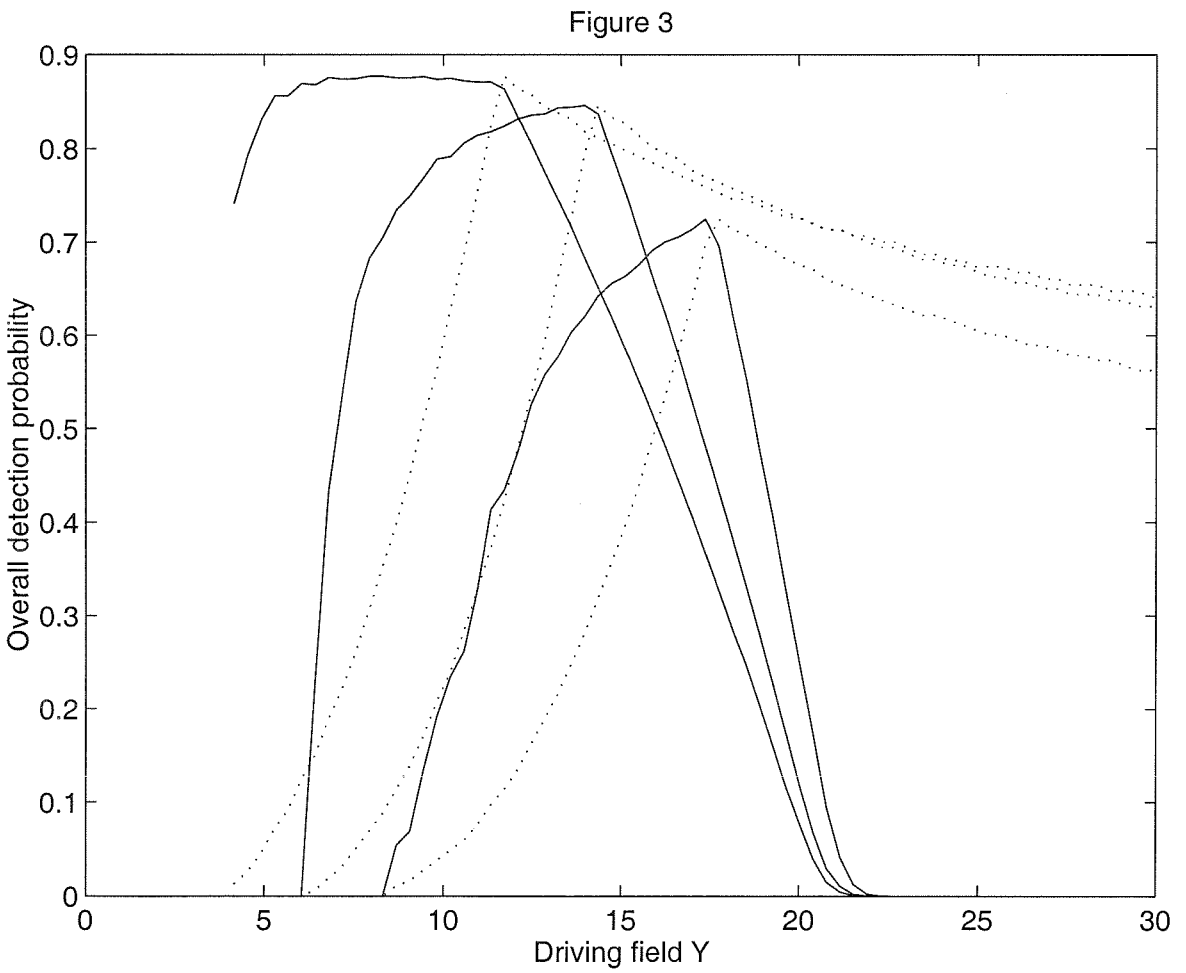

Figure A.3: Solid curves show the overall detection probability versus driving field, for threshold C_0 values of 1, 5, and 10. Dotted lines show the relative size of the test-condition constant c'.

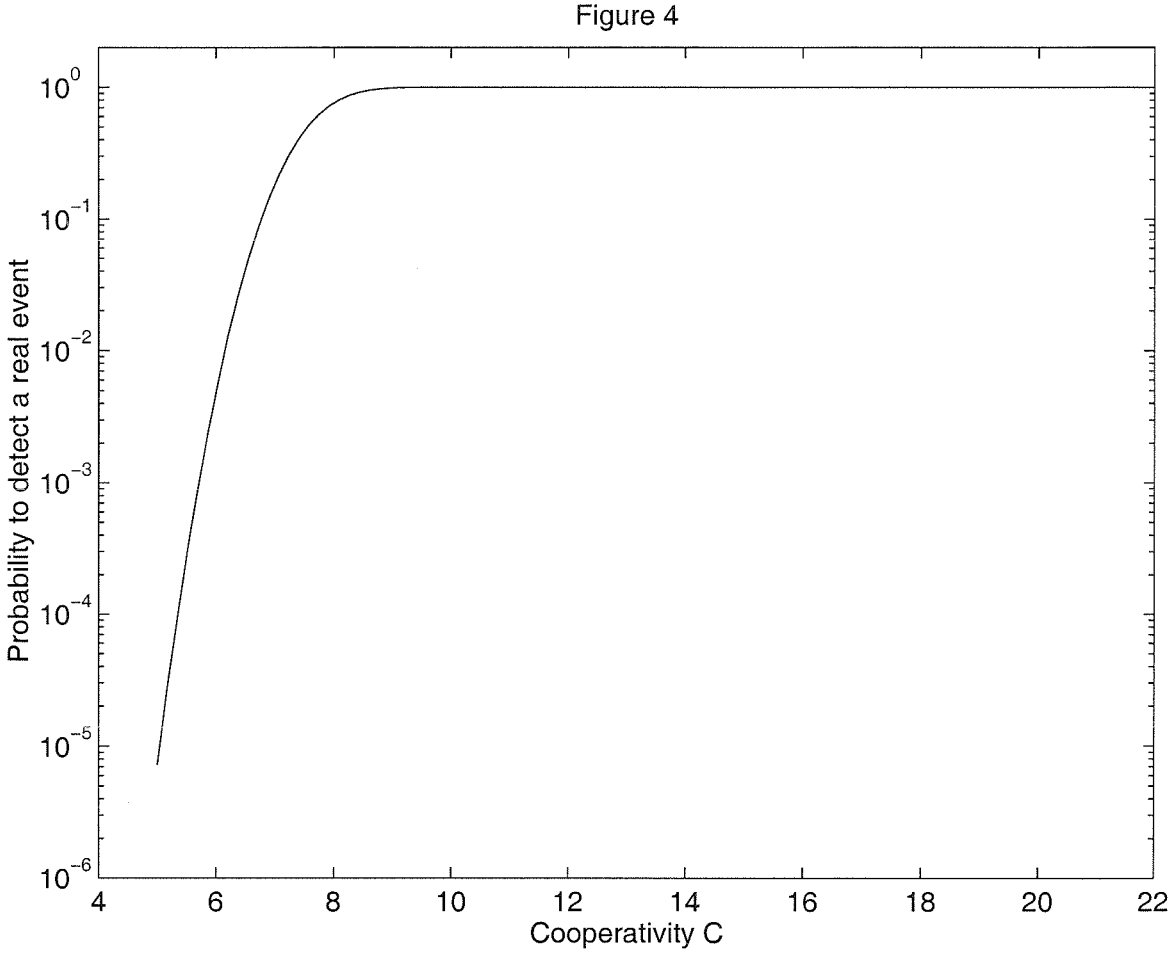

Figure A.4: Power function $\pi(C)$ for $C_0 = 5$.

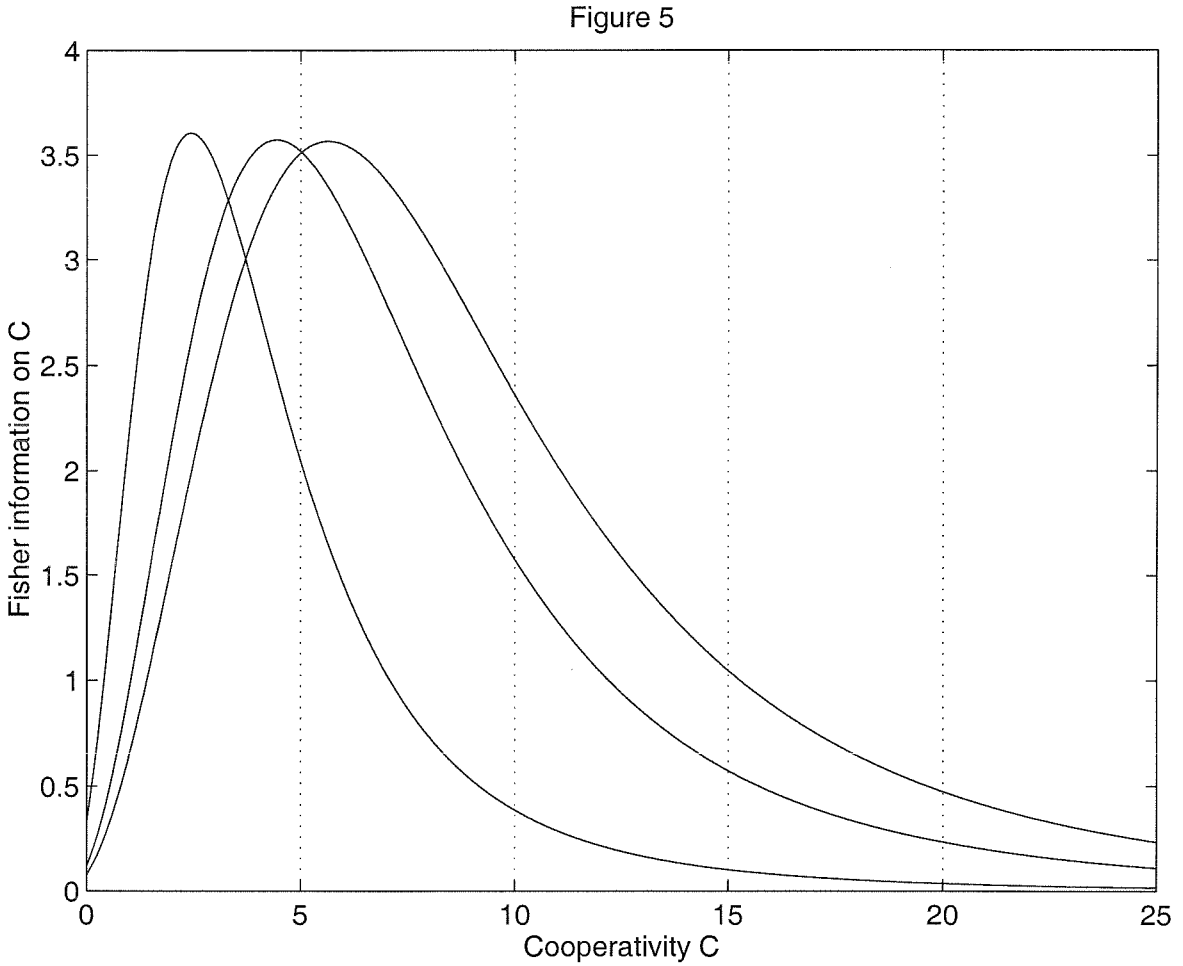

Figure A.5: Fisher information on C versus C, for the optimal Y-values indicated in Figure A.4.

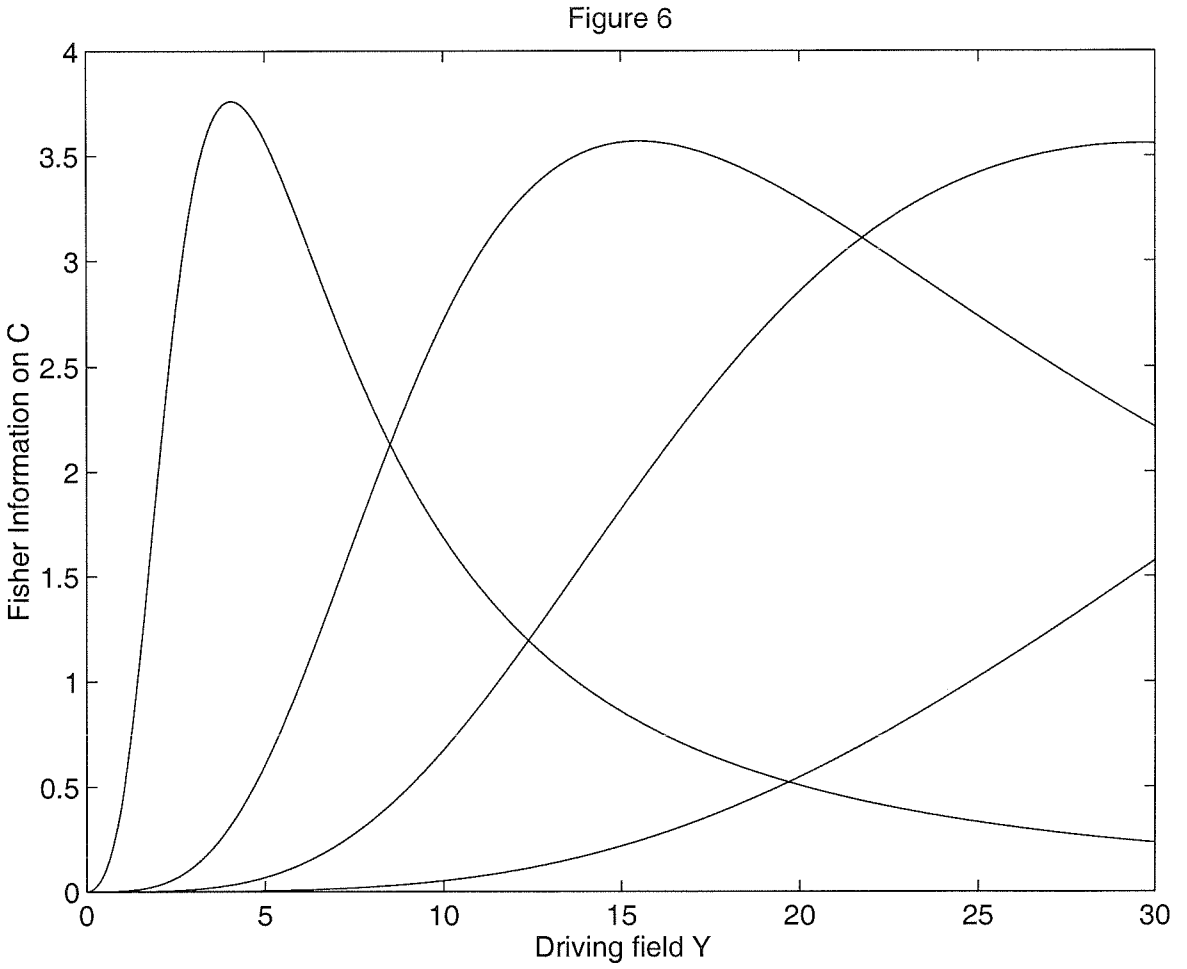

Figure A.6: Fisher information on C versus Y, for C=1, 5, 10, and 22.

Appendix B Quantum feedback in state space

The purpose of this set of notes is to lay out a strategy for exploratory work on feedback control of open quantum systems, with an emphasis on the development of *design principles* for the solution of concrete problems in real experiments. Two particular cavity-QED applications I have in mind are using feedback to trap and cool a single atom inside a cavity, and feedback-stabilized synthesis of arbitrary quantum states of the electromagnetic field. A second principle which guides our approach—in addition to the desire for "practical usability"—is to formulate the control problem in such a way as to make clear precisely which aspects are quantum-mechanical in nature. This speaks for a very explicit treatment of measurement processes, which in turn suggests that we should proceed by building upon *state-space methods* from classical control theory, as opposed to treatments based on the analysis of transfer functions. Such an approach should as well make it easy to include effects associated with measurement backaction.

At the classical (macroscopic) level, feedback is necessary to compensate for the effects of unpredictable disturbances on a system under control, or indeed to make control possible when we simply do not know the initial state of the system. Generally speaking, the use of feedback enables us to control *open* dynamical systems, *i.e.*, systems that are interacting with their environment, and about which we have only incomplete or partly-reliable prior information. To overcome such a handicap we must continually improve our estimate of the evolving system state through real-time measurements (generally also imperfect) made on the system *as* we are trying to control it. Measurement *backaction* greatly complicates the notion of feedback control when the dynamics of the controlled system are manifestly quantum-mechanical. Due to the uncertainty principle, it is rigorously impossible to gain information about

the state of a quantum dynamical system without perturbing it in a manner that cannot be determined beforehand. Hence, any strategy for continually reducing one's uncertainty about the state of an evolving quantum system necessarily induces an excess unpredictable component in its dynamics.

For a certain class of quantum dynamical systems – namely Markovian open quantum systems [9] – stochastic equations can in fact be derived that self-consistently incorporate the effects of continous, indirect observation on the evolution of the system state [8]. Significantly, the form of these equations is such that they may be used to analyze the effects of quantum feedback, which is defined as the process of continuously adjusting the system evolution operator in a manner determined by some function of the measurement results. Having this essential tool for analysis in hand, the next step toward experimental investigations must be to elucidate useful principles for the design of QFC systems. Fortunately, one can draw a very close analogy between open quantum systems and the canonical "plant" of classical control theory making it possible to generalize known principles of control-system design to include Schrödinger-type dynamics as well as quantum-mechanical measurement backaction.

Let's start by having a look at the schematic diagrams shown in Figures 1a and 1b. The *flow diagram* for a classical open-loop plant is traditionally drawn as shown in Figure 1a. The dynamics of the plant are summarized by a differential equation of evolution

$$\dot{\mathbf{x}}(t) = f\left[\mathbf{x}(t), \mathbf{u}(t), \eta_1(t)\right], \tag{B.1}$$

where $\mathbf{x}(t)$ represents all variables required to completely represent the *system state* at time t and $\mathbf{u}(t)$ represents the set of *control inputs* applied at time t. Equation (B.1) is known as the *process* equation. In general, the system evolution (*i.e.*, the function f) can depend on *exogenous* or "noise" variables η_1, which are collectively known as *process noise*. For digital control it is convenient to discretize time according to $t_k = k\zeta$, where k is an integral time index and ζ is a sufficiently small real timestep. Assuming the discretization is fine enough, the process equation may be written as a

difference equation,

$$\mathbf{x}(k) = f\left[\mathbf{x}(k-1), \mathbf{u}(k-1), \eta_1(k-1)\right]. \tag{B.2}$$

In the paradigm of *state feedback control*, one assumes that "direct" knowledge of the system state $\mathbf{x}(k)$ is available to (or least assumed by) the controller, so that the measurement stage is never treated explicitly. For our purposes, however, we should consider the paradigm of *output feedback control*, in which the controller's knowledge of the system state must be inferred from the *system outputs* $\mathbf{y}(k)$. The *measurement equation* (for discrete time) is written

$$\mathbf{y}(k) = g\left[\mathbf{x}(k), \eta_2(k); \mathbf{w}(k)\right]. \tag{B.3}$$

Here $\eta_2(k)$ represents the *measurement noise*, and we include a set of *measurement parameters* $\mathbf{w}(k)$ to allow for strategies in which the measurement function $g[\cdot]$ is *probing* or *adaptive*. In the output feedback scenario, a controller must adjust the control inputs $\mathbf{u}(k)$ in order to control the evolution of the system state $\mathbf{x}(k)$ in a specified manner, basing its control decisions on only the system outputs $\mathbf{y}(k)$. It must do so even in the presence of *uncertainties* $\eta_1(k)$, $\eta_2(k)$.

Figure 1b shows a schematic diagram of typical open quantum system encountered in the context of quantum optics. Such a system will generally have both *c-number control inputs* $\mathbf{u}(t)$ (which appear as parameters in the master equation), and *quantum driving fields* $\hat{b}_{in}(t)$. The system state must now be represented as a density operator $\hat{\rho}(t)$, but the system outputs can still be represented as c-numbers $\mathbf{y}(t)$. In an observed open quantum system, the process and measurement equations do not separate cleanly—the system evolution and measurement results are jointly determined by a stochastic Schrödinger equation. Nevertheless we can formulate a well-posed problem called *feedback control of open quantum systems*, in which the task of the controller is to control the evolution of the system state $\hat{\rho}(t)$ in a specified manner, basing its control decisions on only the *classical record* of measurement re-

sults $\{\mathbf{y}(t_n), \mathbf{y}(t_{n-1}), ...\}$. Again we can throw in uncertainties such as measurement noise for added realism, so that the controller must be able to execute its function having only a noise-corrupted classical record at its disposal.

Figure 2a shows the flow diagram for state-space feedback control of a classical plant. The diagram was drawn to emphasize a very important aspect of state-space control design—feedback controllers can be designed as the cascade of a *state observer* module and a *control law* module. According to the principle of *separability* [70], these two stages can be designed almost completely independently! The function of the state observer is to keep a running *state estimate* $\xi(k)$ for the control system, based upon a record of the applied sequence of controls $\{\mathbf{u}(k), \mathbf{u}(k-1), ...\}$ and the observed system outputs $\{\mathbf{y}(k), \mathbf{y}(k-1), ...\}$. The controller always uses this estimated system state to determine the appropriate set of control inputs to apply during the next timestep, according to $\mathbf{u}(k+1) = F[\xi(k)]$. The idea of separability is that one can treat the problem of state estimation independently from that of determining a good feedback law $F[\cdot]$.

I would propose that the principle of separability could serve as the fundamental basis of feedback design for open quantum systems as well, as suggested in Figure 2b. In the remainder of these notes, I sketch an outline of one possible approach to the formalism.

Quantum process equation: We consider a Markovian open quantum system specified by the Lindblad Master equation (the most general local Master equation for evolution of the density matrix)

$$\begin{aligned}
\frac{d\rho(t)}{dt} &= -i\hbar\left[\hat{H}(\mathbf{u}(t)), \rho(t)\right] \\
&\quad + \frac{1}{2}\sum_{j=1}^{d}\left(2\hat{c}_j\rho(t)\hat{c}_j^\dagger - \hat{c}_j^\dagger\hat{c}_j\rho(t) - \rho(t)\hat{c}_j^\dagger\hat{c}_j\right) \\
&\equiv -i\hbar\left[\hat{H}(\mathbf{u}(t)), \rho(t)\right] + \sum_{j=1}^{d}\mathcal{L}_j\rho(t) \qquad \text{(B.4)} \\
&\equiv \mathcal{L}_{\text{tot}}\rho(t). \qquad \text{(B.5)}
\end{aligned}$$

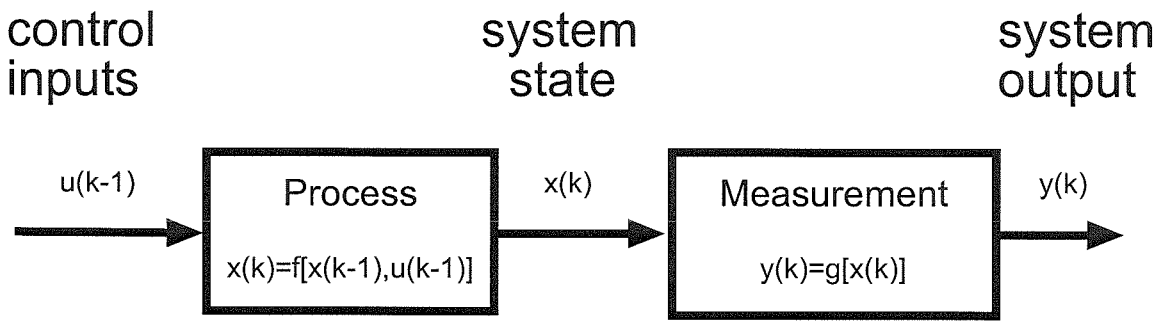

FIGURE 1A

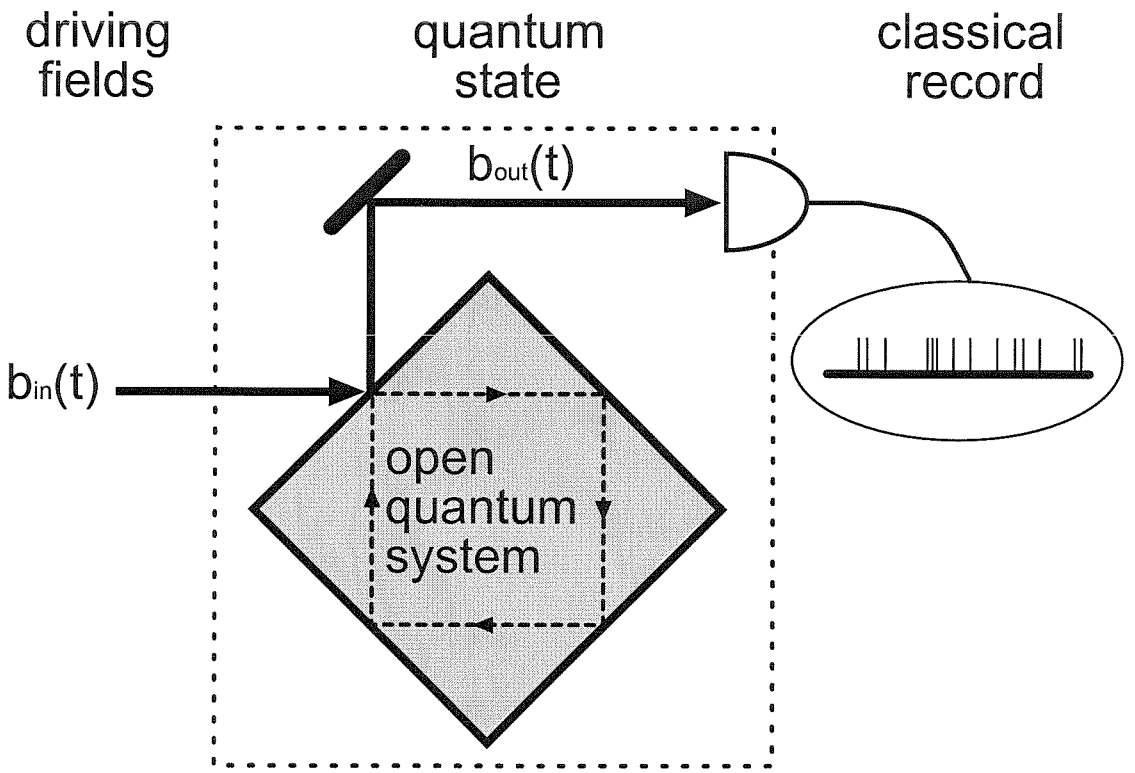

FIGURE 1B

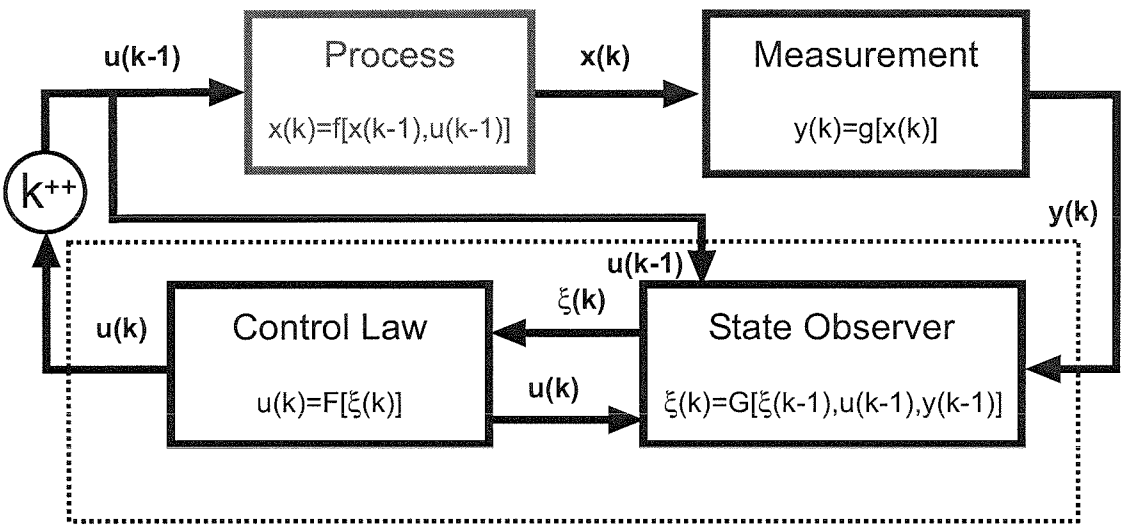

Feedback controller

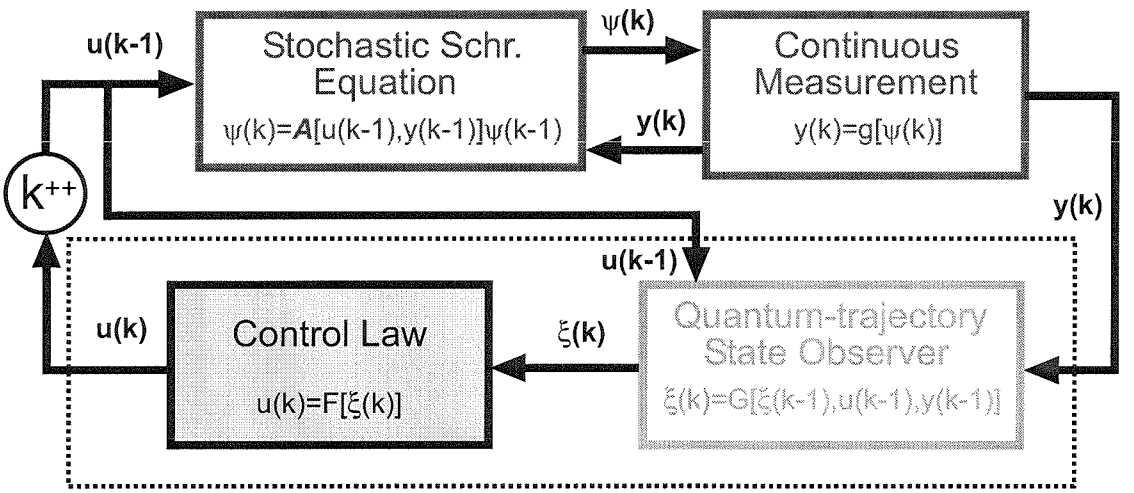

Quantum or Semiclassical feedback controller

Here we write $\hat{H}(\mathbf{u}(t))$ to indicate that the system Hamiltonian depends on a set of parameters $\mathbf{u}$ that may vary with time, and in the second line define superoperators $\mathcal{L}_j$ to represent the partial effect of each dissipative channel on the system evolution. The total evolution superoperator $\mathcal{L}_{\text{tot}}$ is known as the *Liouvillian*.

Simply stated, the central problem in quantum control is to design an autonomous controller that varies $\mathbf{u}(t)$ in order to achieve desired modifications to the system dynamics. Examples of typical control objectives would be to change the steady-state solution of the Master equation (B.4), to modify the eigenvalues of the Liouvillian, or to cause the system to cycle through some desired progression of "target" states. In the feedback control scenario, it is assumed that the controller bases its actions on the outcomes of *measurements* made on the reservoirs associated with selected dissipative channels $\mathcal{L}_j$.

Let us assume that d_{obs} of the d dissipative channels are subject to observation by the controller. In the ideal case of continuous perfect detection, the evolution of the observed system would obey a stochastic master equation such that

$$\rho(t+dt) \to \frac{\hat{c}_j \rho(t) \hat{c}_j^\dagger}{\text{Tr}\left[\hat{c}_j \rho(t) \hat{c}_j^\dagger\right]} \tag{B.6}$$

following a detection event in the j^{th} channel, and

$$\begin{aligned}
\frac{d\rho(t)}{dt} &= -i\hbar\left[\hat{H}(\mathbf{u}(t)), \rho(t)\right] \\
&\quad -\frac{1}{2}\sum_{j=1}^{d_{\text{obs}}}\left(\hat{c}_j^\dagger \hat{c}_j \rho(t) + \rho(t)\hat{c}_j^\dagger \hat{c}_j\right) + \sum_{j=d_{\text{obs}}}^{d} \mathcal{L}_j \rho(t) \\
&\equiv -i\hbar\left[\hat{H}_{\text{eff}}(\mathbf{u}(t)), \rho(t)\right] + \sum_{j=d_{\text{obs}}}^{d} \mathcal{L}_j \rho(t)
\end{aligned} \tag{B.7}$$

during periods of time between detected quantum jumps. This defines the *effective Hamiltonian* $\hat{H}_{\text{eff}}(\mathbf{u}(t)) \equiv \hat{H}(\mathbf{u}(t)) - \frac{i}{2}\sum_{j=1}^{d_{\text{obs}}} \hat{c}_j^\dagger \hat{c}_j$ for the partially observed system.

We next transform these equations to a discrete-time formulation $dt \to \Delta t$ with integer time index $n \equiv t/\Delta t$, retaining terms only to first order in the time increment

Δt. Hence we have the discrete *observed-process equation*

$$\begin{aligned}\rho(n+1) &= \rho(n) - i\hbar\Delta t\left[\hat{H}_{\text{eff}}(\mathbf{u}(n)), \rho(n)\right] \\ &\quad + \sum_{j=d_{\text{obs}}}^{d} \Delta t\, \mathcal{L}_j\, \rho(n), \qquad \text{(no jumps detected)} \\ &= \frac{\hat{c}_j \rho(n)\, \hat{c}_j^\dagger}{\text{Tr}\left[\hat{c}_j \rho(n)\, \hat{c}_j^\dagger\right]}. \qquad \text{(jump detected inchannel } j\text{)} \end{aligned} \tag{B.8}$$

The probability for a quantum jump of type j to occur during a timestep of duration Δt is given by $\Pr(j) = \Delta t\, \text{Tr}\left[\rho(n)\, \hat{c}_j^\dagger \hat{c}_j\right]$.

Inference-update equation: Unfortunately, it will never be possible to detect these quantum jumps with perfect efficiency in any real experimental setup. In general, the detectors monitoring the system under control will register an *ostensible* sequence of quantum jumps $\{m_i\}$ that corresponds to a noisy version of the "true" sequence $\{j_i\}$. If an event m_i is detected at time n_i, we would like to have a rule for updating our conditional density matrix (estimate of the evolving system state) that embodies some sort of optimal filtration of the detection noise. A straightforward application of Bayes' rule yields

$$\Pr(j_i|m_i, \tilde{\rho}_i) = \frac{\Pr(m_i|j_i)\Pr(j_i)}{\Pr(m_i)}, \tag{B.9}$$

where the "prior distribution" $\Pr(j_i) = \text{Tr}[\hat{c}_{j_i}^\dagger \hat{c}_{j_i} \tilde{\rho}_i]$, and $\Pr(m_i|j_i)$ is given by the description of the detector noise/inefficiency. Note that $\tilde{\rho}_i$ with a tilde represents the controller's *internal* (recursive) *estimate* of the evolving system state, and we have introduced the simplifying notation $\tilde{\rho}(n_i) \leftrightarrow \tilde{\rho}_i$. Averaging over quantum jumps,

$$\tilde{\rho}_{i+1} = \sum_{j_i} \Pr(j_i \,|\, m_i)\left[\hat{c}_{j_i} \tilde{\rho}_i \hat{c}_{j_i}^\dagger \,/\, ||\cdot||\right] \tag{B.10}$$

$$= \frac{1}{\Pr(m_i)} \sum_{j_i} \Pr(m_i \,|\, j_i)\Pr(j_i)\left[\hat{c}_{j_i} \tilde{\rho}_i \hat{c}_{j_i}^\dagger \,/\, ||\cdot||\right] \tag{B.11}$$

$$= \frac{1}{\Pr(m_i)} \sum_{j_i} \Pr(m_i \,|\, j_i)\, \hat{c}_{j_i} \tilde{\rho}_i \hat{c}_{j_i}^\dagger\,. \tag{B.12}$$

Recall that m_i is here considered to be an observed, fixed value. Also, note that the marginal probability for observing the detector response m_i is given by the trace of $\sum_{j_i} \Pr(m_i \mid j_i)\ \hat{c}_{j_i} \rho_i \hat{c}_{j_i}^{\dagger}$, so ρ_{i+1} is indeed normalized.

The form of the control law for determining $\mathbf{u}(n)$ thus takes the form $\mathbf{u}_i = \mathcal{F}[\tilde{\rho}_i]$, where $\mathcal{F}$ may *a priori* be some arbitrarily complicated functional. Once a fixed control law has been designed, however, the task of simulating the *closed loop* system evolution reduces to jointly evolving the coupled stochastic Master equation (B.4) and the recursive update rule (B.12).

Optimal control equation: In this section we formulate recurrence equations for optimal control design, relative to a given *cost function* $\mathcal{R}[\rho(n), \mathbf{u}(n), n]$. We assume that $\mathcal{R}$ is of the form

$$\mathcal{R}[\rho(n), \mathbf{u}(n), n] = \sum_a P_t(a)\ \mathrm{Tr}\left[\hat{E}_t^a \rho(n)\right] + Q_t[\mathbf{u}(n)],$$

where the set $\left\{\hat{E}_t^a\right\}$ defines an aribtrary postive operator-valued measure (POVM) on the system state space. Here $P_t(a)$ is any real-valued function of the POVM index a, and Q_t is any real-valued function of the control parameters $\mathbf{u}(n)$. By restricting $\mathcal{R}[\rho(n), \mathbf{u}(n), n]$ to depend on $\rho(n)$ only through the outcome of a POVM, we achieve linearity with respect to probabilistic ensembles of system states while still allowing the control problem to be specified in terms of any *measurable* property of the system state at time n.

The objective of optimal control is to determine the control law that minimizes the *expected* value of the total cost $I_N[\rho(1)]$,

$$I_N[\rho(1)] \equiv \sum_{n=1}^{N} \mathcal{R}[\rho(n), \mathbf{u}(n), n]. \tag{B.13}$$

In all but simple cases restricted to linear instantaneous feedback, it is not yet known how to derive a closed analytical form for this equation. However, the problem can certainly be specified in terms of coupled difference equations amenable to numerical solution. We therefore propose to undertake the development of efficient algorithms

to solve such optimization problems for a broad class of quantum-optical open systems that will be the subject of experimental investigations in the near future.

Bibliography

[1] K. Kraus, *States, effects, and operations: fundamental notions of quantum theory* (Springer, Berlin 1983).

[2] J. Preskill, lecture notes for Caltech Ph229, available (as of 4/98) on the World Wide Web at *http://www.theory.caltech.edu/~preskill/ph229.*

[3] M. A. Nielsen, C. M. Caves, B. Schumacher, and H. Barnum, "Information-theoretic approach to quantum error-correction and reversible measurement," Proc. Roy. Soc. London A **454**, 277-304 (1998).

[4] H. M. Wiseman, "Quantum trajectories and quantum measurement theory," Quantum Semiclass. Opt. **8**, 205-222 (1996).

[5] V. B. Braginsky and F. Ya. Khalili, *Quantum Measurement* (Cambridge University Press, Cambridge, 1992).

[6] C. A. Fuchs, *Distinguishability and Accessible Information in Quantum Theory*, University of New Mexico Doctoral Dissertation (1995).

[7] H. M. Wiseman and G. J. Milburn, "Quantum theory of field-quadrature measurements," Phys. Rev. A **47**, 642-662 (1993).

[8] H. J. Carmichael, *An Open Systems Approach to Quantum Optics* (Springer, Berlin 1993).

[9] C. W. Gardiner, *Quantum Noise* (Springer, Berlin 1991).

[10] H. Mabuchi and P. Zoller, "Inversion of Quantum Jumps in Quantum-Optical Systems Under Continuous Observation," Phys. Rev. Lett. **76**, 3108-3111 (1996).

[11] M. A. Nielsen and C. M. Caves, "Reversible quantum operations and their application to teleportation," Phys. Rev. A **55**, 2547-2556 (1997).

[12] C. A. Fuchs and A. Peres, "Quantum-state disturbance versus information gain – uncertainty relations for quantum information," Phys. Rev. A **53**, 2038-2045 (1996).

[13] P. Zoller and C. W. Gardiner, "Quantum noise in quantum optics: the Stochastic Schrödinger equation," in *Quantum fluctuations, Session LXIII of l'Ecole d'Ete de Physique des Houches* (North-Holland, Amsterdam 1997).

[14] H. Mabuchi, "Dynamical identification of open quantum systems," Quantum Semiclass. Opt. **8**, 1103-1108 (1996).

[15] H. M. Wiseman, "Adaptive phase measurements of optical modes – going beyond the marginal Q-distribution," Phys. Rev. Lett. **75**, 4587-4590 (1995).

[16] A. C. Doherty, A. S. Parkins, S. M. Tan, and D. F. Walls, "Motional states of atoms in cavity QED," to appear in Phys. Rev. A (tentatively scheduled for June, 1998).

[17] J. A. Dunningham, H. M. Wiseman, and D. F. Walls, "Manipulating the motion of a single atom in a standing wave via feedback," Phys. Rev. A **55**, 1398-1411 (1997).

[18] Q. A. Turchette, *Quantum optics with single atoms and single photons*, California Institute of Technology doctoral dissertation (1997).

[19] R. J. Thompson, *Cavity quantum electrodynamics in the optical domain: structure and dynamics in the strong coupling regime*, University of Texas doctoral dissertation (1994).

[20] K. An, J. J. Childs, R. R. Dasari, and M. S. Feld, "Microlaser – a laser with one-atom in an optical-resonator," Phys. Rev. Lett. **73**, 3375-3378 (1994).

[21] Y. F. Zhu, D. J. Gauthier, S. E. Morin, Q. L. Wu, H. J. Carmichael, and T. W. Mossberg, "Vaccum Rabi splitting as a feature of linear-dispersion theory – analysis and experimental-observation," Phys. Rev. Lett. **64**, 2499-2502 (1990).

[22] H. J. Kimble: in *Cavity Quantum Electrodynamics*, edited by P. Berman (Academic Press, San Diego 1994).

[23] R. J. Thompson, Q. Turchette, O. Carnal, and H. J. Kimble, "Nonlinear spectroscopy in the strong-coupling regime of cavity QED," Phys. Rev. A **57**, 3084-3104 (1998).

[24] Q. A. Turchette, R. J. Thompson, and H. J. Kimble, "One-dimensional atoms," Appl. Phys. B **60**, S1-S10 (1995).

[25] Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble, "Measurement of conditional phase shifts for quantum logic," Phys. Rev. Lett. **75**, 4710-4713 (1995).

[26] 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**, 1393-1395 (1996).

[27] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, "Trapping of neutral sodium atoms with radiation pressure," Phys. Rev. Lett. **59**, 2631-2634 (1987).

[28] C. M. Caves, "Defense of the standard quantum limit for free-mass position," Phys. Rev. Lett. **54**, 2465-2468 (1985).

[29] A. C. Doherty, A. S. Parkins, S. M. Tan, and D. F. Walls, "Motion of a 2-level atom in an optical cavity," Phys. Rev. A **56**, 833-844 (1997).

[30] R. Quadt, M. Collett, and D. F. Walls, "Measurement of atomic motion in a standing light-field by homodyne detection," Phys. Rev. Lett. **74**, 351-354 (1995).

[31] W. E. Moerner, "Examining nanoenvironments in solids on the scale of a single, isolated impurity molecule," Science **265**, 46-53 (1994).

[32] B. C. Stipe, M. A. Rezaei, and W. Ho, "Inducing and viewing the rotational motion of a single molecule," Science **279**, 1907-1909 (1998).

[33] D. M. Meekhof, C. Monroe, B. E. King, W. M. Itano, and D. J. Wineland, "Generation of nonclassical motional states of a trapped atom," Phys. Rev. Lett. **76**, 1796-1799 (1996).

[34] D. Leibfried, D. M. Meekhof, B. E. King, C. Monroe, W. M. Itano, and D. J. Wineland, "Experimental-determination of the motional quantum state of a trapped atom," Phys. Rev. Lett. **77**, 4281-4285 (1996).

[35] M. R. Andrews, D. M. Kurn, H. J. Miesner, D. S. Durfee, C. G. Townsend, S. Inouye, and W. Ketterle, "Propagation of sound in a Bose-Einstein condensate," Phys. Rev. Lett. **79**, 553 (1997).

[36] M. Kozuma, K. Nakagawa, W. Jhe, and M. Ohtsu, "Phase-change of a probe due to oscillation of cold atoms in an optical standing-wave," Phys. Rev. A **57**, R24-R27 (1998); M. Kozuma, Y. Imai, K. Nakagawa, and M. Ohtsu, "Observation of a transient-response of recoil-induced resonance – a method for the measurement of atomic motion in an optical standing-wave," Phys. Rev. A **52**, R3421-R3424 (1995).

[37] S. Guibal, C. Triche, L. Guidoni, P. Verkerk, and G. Grynberg, "Recoil-induced resonances of cesium atoms in the transient domain," Opt. Commun. **131**, 61-65 (1996).

[38] C. M. Caves and G. J. Milburn, "Quantum-mechanical model for continuous position measurements," Phys. Rev. A **36**, 5543-5555 (1987).

[39] H. M. Wiseman and G. J. Milburn, "Interpretation of quantum jump and diffusion-processes illustrated on the Bloch Sphere," Phys. Rev. A **47**, 1652-1666 (1993).

[40] H. M. Wiseman, "Quantum-theory of continuous feedback," Phys. Rev. A **49**, 2133-2150 (1994).

[41] R. Alicki, D. Makowiec, and W. Miklaszewski, "Quantum chaos in terms of entropy for a periodically kicked top," Phys. Rev. Lett. **77**, 838-841 (1996).

[42] R. Schack and C. M. Caves, "Information-theoretic characterization of quantum chaos," Phys. Rev. E **53**, 3257-3270 (1996).

[43] A. H. Zewail, "Femtochemistry – recent progress in studies of dynamics and control of reactions and their transition-states," J. Phys. Chem. **100**, 12701-12724 (1996).

[44] S. D. Gensemer and P. L. Gould, "Ultracold collisions observed in real-time," Phys. Rev. Lett. **80**, 936-939 (1998).

[45] C. J. Hood, M. S. Chapman, T. W. Lynn, and H. J. Kimble, "Real-time cavity QED with single atoms," to appear in Phys. Rev. Lett. (1998).

[46] M. J. Collett, R. Loudon, and C. W. Gardiner, "Quantum-theory of optical homodyne and heterodyne-detection," J. Mod. Opt. **34**, 881-902 (1987).

[47] L. A. Lugiato, "Theory of optical bistability," Prog. Opt. **21**, 71-216 (1984).

[48] P. Marte, R. Dum, R. Taïeb, and P. Zoller, "Resonance fluorescence from quantized one-dimensional molasses," Phys. Rev. A **47**, 1378-1390 (1993).

[49] O. Schmidt, K.-M. Knaak, R. Wynands, and D. Meschede, "Cesium saturation spectroscopy revisited: How to reverse peaks and observe narrow resonances," Appl. Phys. B **59**, 167-178 (1994).

[50] Superpolished BK7 substrates with multi-layer dielectric coatings fabricated by Research Electrooptics in Boulder, Colorado.

[51] A. E. Siegman, *Lasers* (University Science Books, Mill Valley 1986).

[52] R. J. Rafac, C. E. Tanner, A. E. Livingston, K. W. Kukla, H. G. Berry, and C. A. Kurtz, "Precision lifetime measurements of the $6P_{1/2,3/2}$ states in atomic cesium," Phys. Rev. A **50**, R1976-R1979 (1994).

[53] J. H. Shirley, "Modulation transfer processes in optical heterodyne saturation spectroscopy," Opt. Lett. **7** 537-539 (1982).

[54] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, "Laser phase and frequency stabilization using an optical-resonator," Appl. Phys. B **31**, 97-105 (1983).

[55] T. L. Boyd and H. J. Kimble, "Frequency stabilization of a continuous-wave Ti:sapphire laser," Opt. Lett. **16**, 808-810 (1991).

[56] J. Ye, *Ultrasensitive high resolution laser spectroscopy and its application to optical frequency standards*, University of Colorado doctoral dissertation (1997).

[57] A. Cable, M. Prentiss, and N. P. Bigelow, "Observations of sodium atoms in a magnetic molasses trap loaded by a continuous uncooled source," Opt. Lett. **15**, 507-509 (1990).

[58] E. S. Polzik, J. Carri, and H. J. Kimble, "Atomic spectroscopy with squeezed light for sensitivity beyond the vacuum-state limit," Appl. Phys. B **55**, 279-290 (1992).

[59] G. Rempe, R. J. Thompson, and H. J. Kimble, "Measurement of ultralow losses in an optical interferometer," Opt. Lett. **17**, 363-365 (1992).

[60] D. R. T. Appadoo, P. F. Bernath, and R. J. Le Roy, "High-resolution visible spectrum for the A $^3\Pi(1)$-$X^1\Sigma^+$ system of IBr," Can. J. Phys. **72**, 1265- 1272 (1994).

[61] W. G. Brown, "Absorption spectrum of Iodine Bromide," Phys. Rev. **42**, 355-363 (1932).

[62] K. G. Libbrecht and J. L. Hall, "A low-noise high-speed diode-laser current controller," Rev. Sci. Instr. **64**, 2133-2135 (1993).

[63] H. P. Yuen, "Contractive states and the standard quantum limit for monitoring free-mass positions," Phys. Rev. Lett. **51**, 719-722 (1983).

[64] M. Ozawa, "Measurement breaking the standard quantum limit for free-mass position," Phys. Rev. Lett. **60**, 385-388 (1988).

[65] M. T. Jaekel and S. Reynaud, "Quantum limits in interferometric measurements," Europhys. Lett. **13**, 301-306 (1990).

[66] V. B. Braginsky and Yu. I. Vorontsov, Ups. Fiz. Nauk **114**, 41 (1974) [Sov. Phys. Usp. **17**, 644 (1975)].

[67] C. M. Caves, K. S. Thorne, R. W. P. Drever, V. D. Sandberg, and M. Zimmermann, Rev. Mod. Phys. **52**, 341 (1980).

[68] G. J. Milburn, K. Jacobs, and D. F. Walls, "Quantum-limited measurements with the atomic-force microscope," Phys. Rev. A **50**, 5256-5263 (1994).

[69] H. Mabuchi, J. Ye, and H. J. Kimble, preprint (1998).

[70] O. L. R. Jacobs, *Introduction to Control Theory* (Oxford University Press, Oxford, 1993).

[71] G. J. Milburn, "Classical and quantum conditional statistical dynamics," Quantum Semiclass. Opt. **8**, 269-276 (1996).

[72] V. P. Belavkin and P. Staszewski, "Nondemolition observation of a free quantum particle," Phys. Rev. A **45**, 1347-1356 (1992).

[73] C. W. Gardiner, *Handbook of Stochastic Processes* (Springer, Berlin 1990).

[74] Equations (4.16)-(4.17) can be derived by direct solution of the Schrödinger equation with a Gaussian ansatz.

[75] C. W. Helstrom, *Elements of Signal Detection and Estimation* (PTR Prentice-Hall, Englewood Cliffs 1995).

[76] S. L. Braunstein, A. S. Lane, and C. M. Caves, "Maximum-likelihood analysis of multiple quantum phase measurements," Phys. Rev. Lett. **69**, 2153-2156 (1992).

[77] B. C. Sanders and G. J. Milburn, "Optimal quantum measurements for phase estimation," Phys. Rev. Lett. **75**, 2944-2947 (1995).

[78] S. L. Braunstein and C. M. Caves, "Statistical distance and the geometry of quantum states," Phys. Rev. Lett. **72**, 3439-3442 (1994).

[79] A. S. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (North-Holland, Amsterdam 1982).

[80] Recent work in quantum optics has investigated *tomographic* reconstruction of density matrices, the procedure for which involves making sets of related but non-identical measurements. For example, see:
U. Leonhardt, H. Paul, and G. M. Dariano, "Tomographic reconstruction of the density-matrix via pattern functions," Phys. Rev. A **52**, 4899-4907 (1995) (and references therein); J. F. Poyatos, R. Walser, J.-I. Cirac, P. Zoller, and R. Blatt, "Motion tomography of a single trapped ion," Phys. Rev. A **53**, R1966-R1969 (1996).

[81] L. Ljung, *System Identification: Theory for the User* (PTR Prentice Hall, Upper Saddle River 1987).

[82] Wiseman and Milburn (1993) have derived closed-loop master equations to describe the effects of feeding an observed output channel back to a controlled system input.
H. M. Wiseman and G. J. Milburn, "Quantum-theory of optical feedback via homodyne detection," Phys. Rev. Lett. **70**, 548-551 (1993).

[83] M. H. DeGroot, *Probability and Statistics, second edition* (Addison-Wesley, Reading 1989).

[84] C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and P. Gould, "Resolved-side-band Raman cooling of a bound atom to the 3D zero-point energy," Phys. Rev. Lett. **75**, 4011-4014 (1995).

[85] C. M. Caves, "Quantum-mechanics of measurements distributed in time – a path-integral formulation," Phys. Rev. D **33**, 1643-1665 (1986).

[86] Although an investigation of the associated measurement *back-action* would be quite interesting, such considerations lie beyond the scope of the current work.

[87] In fact this scenario is not so far from experimental reality, since avalanche photodiodes can provide quantum efficiencies $\geq 70\%$ in the near infrared. The problem of having to collect atomic fluorescence into 4π solid angle can be avoided by using a cavity for which $\kappa \gg g^2/\kappa > \gamma_\perp$, corresponding to the *bad cavity regime* investigated theoretically in [93] and experimentally in [24]. The main technical obstacle to implementing the QPE procedure as described above would most likely be the dead-time of available photon-counting modules, which currently amounts to tens of nanoseconds.

[88] The simulations were performed using the *Quantum Optics Toolbox for Matlab* written by S. M. Tan (private communicaton).

[89] 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 (1992).

[90] 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**, 1727-1730 (1991).

[91] P. Alsing and H. J. Carmichael, "Spontaneous dressed-state polarization of a coupled atom and cavity mode," Quantum Opt. **3**, 13-32 (1991).

[92] Indeed, one would like to quantify the information rate associated with the observation of classical records for different driving conditions. This would entail computing the Fisher information on $\vec{\theta}$ contained in the stochastic ensemble of records Ξ. Note that for a fixed effective Hamiltonian, the Fisher information rate on a given parameter could be optimized by varying the basis in which the output channels are measured.

[93] P. Rice and H. J. Carmichael, "Single-atom cavity-enhanced absorption: 1. Photon statistics in the bad-cavity limit," IEEE J. Quantum Electron. **24**, 1351-1366 (1988).