Chapter 1

Introduction and Motivation

1.1 Introduction

Detecting and directing the evolution of simple quantum systems is a growing field of study which addresses fundamental aspects of quantum theory as well as emerging technological goals in quantum information science. A handful of physical systems possess the characteristics that allow experiments in this field: coherent interaction between simple quantum systems, comparably weak dissipative couplings to the environment, and accessible input and output channels for control and detection of the dynamics. One such system is cavity quantum electrodynamics, in which a single atom interacts strongly with the quantized field of a high-finesse optical cavity.

The coherent exchange of excitation between atomic dipole and cavity field defines the quantum system under consideration. This system couples to the environment via atomic spontaneous emission and also through the decay of the cavity field. A large part of this cavity decay is associated with transmission of light through the cavity mirrors, which leads to a well-defined output channel where information can in fact be "rescued" through appropriate detection or use of the cavity output light. Likewise, the opportunity to couple light *into* the cavity mode with arbitrary power and frequency provides a convenient input channel for real-time manipulation of the atom-cavity state. When the coherent coupling dominates the decay rates, we have a system in the *strong coupling* regime which has distinctly nonclassical behavior despite its interactions with the environment.

1.2 Motivations

Cavity QED in the strong coupling regime offers the possibility for efficient measurement and control of single quanta, and for rapid and controlled coherent interactions between these single quanta. In the language of quantum information theory, cavity QED is one of several viable platforms for quantum logic and quantum communication. Ongoing technical progress brings the field closer to achieving atomic position control that is fine enough and stable enough to perform a series of atom-field logic gates at high fidelity; position control is required for this purpose because it means a precise knowledge of the coherent coupling rate $g(\vec{r})$. This ability in turn will allow for on-demand atom-cavity interactions to prepare and coherently couple novel quantum states of the atom and field.

For purposes of quantum information science, atomic physics in optical cavities has the advantage of offering clock rates that are fast in absolute terms, with coherent coupling rates for current experiments in the range $g_0/2\pi \approx 100 - 200$ MHz. Its chief strength, however, may lie in the marriage of atomic internal states, easily accessible for preparation and robust enough for storage, with states of the light field which can be easily and rapidly transported across large distances. In other words, optical cavity QED provides an attractive setting for the implementation of diverse protocols in quantum communication, quantum teleportation and entanglement distribution, and thus eventually extended quantum networking.

Precise position measurement and control are important to allow high-fidelity quantum gates but also in their own right for what these attempts can teach us about how to measure and steer a quantum system with the handles we are given. Quantum state estimation and control, as well as the implementation of "designer" evolution schemes using active feedback, form an exciting area of quantum information science today.

1.3 Introduction to the Cavity QED System

Optical elements in general present boundary conditions that alter the free-space quantization structure of the electromagnetic field. This modified electromagnetic mode structure in turn affects the interactions of an atomic dipole with light, including decay into the now-altered vacuum. Diverse observations have demonstrated changes in atomic radiative processes caused by the presence of a boundary; for example, boundary-induced atomic level shifts form the basis of the Casimir effect and numerous other phenomena [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16]. For a review of these effects and their development into the field of cavity QED, see for example Ref. [17].

The boundary conditions imposed by an optical cavity in particular create a set of electromagnetic field modes confined between the cavity mirrors. These resonant cavity modes are well defined in frequency and in spatial structure. The modes of an optical cavity typically subtend a small fraction of 4π in solid angle, and thus do not significantly suppress free-space atomic spontaneous emission. However, the presence of the cavity introduces a new rate, the rate of coherent exchange of excitation between atom and cavity field. Through this coherent coupling the atom and cavity decay linewidths do in fact alter one another, at first perturbatively and then strongly as the coherent coupling becomes large relative to both decays. Finally, when the physical size of a cavity is reduced until the cavity mode volume is near the atomic "radiative" volume, a whole new set of quantum dynamics associated with the full quantum susceptibility can be explored within the setting of cavity QED.

Figure 1.1 depicts the components and rates essential for understanding optical cavity QED. A Fabry-Perot resonator is created by aligning two highly reflective spherical mirrors at separation l measured along the cavity axis. Modes of the cavity possess a standing-wave structure along the axis, so the cavity supports a set of longitudinal modes separated in frequency by a free spectral range (FSR) of $c/2l$ where c is the speed of light. (Mirror coatings cause the FSR to deviate very slightly from this simple formula, as discussed in [18] and briefly in Chapters 6 and 7.) At

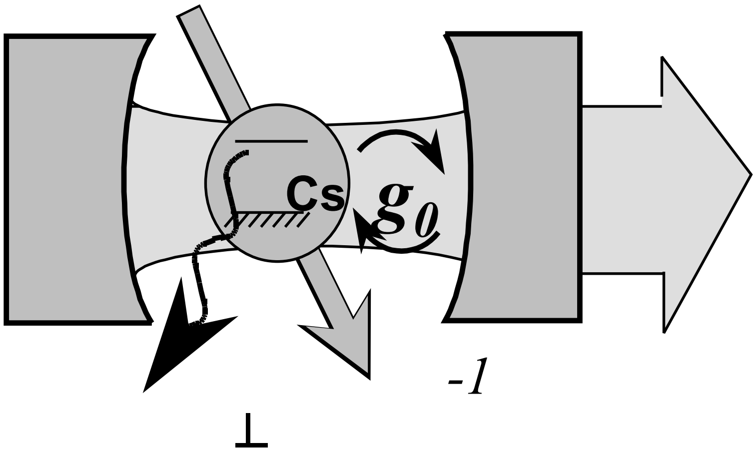

Figure 1.1: Basic rates in the cavity QED system

each longitudinal mode the cavity supports a complete set of transverse modes of different transverse spatial profiles. The TEM_{00} mode has a cylindrically symmetric Gaussian profile, characterized by a beam waist w_0 for the cavity field.

If the two mirrors are assumed to be identical, the cavity is characterized by the per-mirror power transmission T and loss A. The total empty cavity round-trip loss is $L_{cav} = 2(\mathrm{T} + \mathrm{A})$. The cavity finesse ($F$) is given by

$$F = \frac{2\pi}{L_{cav}} = \frac{\pi}{\mathrm{T} + \mathrm{A}}. \tag{1.1}$$

The finesse can also be expressed as the ratio of free spectral range to cavity linewidth. It is closely related to another commonly used quantity, the resonator quality factor Q, which is the ratio of the resonant optical frequency to the cavity decay linewidth. The use of finesse is attractive as it depends only on the mirror properties and not strongly on the cavity dimension. When the cavity length changes, both the FSR and linewidth of the cavity vary as the inverse of the cavity length, and hence the finesse remains a constant. Q and F are related by the ratio of the optical frequency to the FSR.

When a single atom is present in the cavity mode volume, the atomic dipole interacts with the electric field built up in the cavity mode to provide a coherent coupling between atom and cavity field. That coherent interaction is characterized by the rate g_0, which is one half the Rabi frequency for a ground state atom to couple to a single photon in the cavity field. Other important rates in the system are those which characterize decay. These include the atomic spontaneous emission, $\gamma_\perp$, and the cavity decay rate κ. To correspond with the use of g_0 for the coherent interaction rate, cavity QED literature typically quotes both atomic and cavity decay in field decay rates, i.e., half the full "power" decay rates. Thus $\kappa/2\pi = \frac{1}{2}\frac{FSR}{Finesse}$, and $\gamma_\perp = 1/2\tau_0$ where τ_0 is the atomic radiative lifetime. A final important quantity is the rate for a single atom to traverse the cavity field and thus to move from zero to optimal coupling. If the relevant motional timescale is given by a transit or orbit time τ, then $1/\tau$ can be considered a third decay rate characterizing the atom-cavity interaction. Strong coupling for both internal and external degrees of freedom is ensured when $g_0 \gg (\gamma_\perp, \kappa, 1/\tau)$. The cavity QED experiments I have been involved in all take place in this limit, where the atom and cavity jointly form a quantum system open to the environment through decay rates small compared to their coherent exchange of excitation.

1.4 History of Us

My involvement in cavity QED research as a student has encompassed three experimental setups for cold-atom cavity QED with very short cavities. When I joined the lab in fall of 1996, Christina Hood and Mike Chapman were just building up the second cold atoms experiment in the group, and the first to use very short ($l \approx 10\mu m$) cavities. I was privileged to be involved in the completion of that setup and in the experiment which followed. That experiment is presented in Christina Hood's thesis [19] and in Ref. [20], but I mention here two principal results from that work: first, the demonstration of quantum rather than semiclassical saturation behavior for the atom-cavity system, and second, the mapping out of the vacuum Rabi splitting – measured atom by atom – which clearly demonstrated mechanical effects of the cavity QED probe on single atoms transiting the mode.

At the conclusion of the vacuum Rabi splitting experiment, inspired partially by the degradation of our physics cavity finesse, Christina and I embarked on a rebuilding of the experiment with an improved cavity. This cavity, which became the heart of the "atom-cavity microscope" experiment, had a finesse of 480,000 as opposed to the 180,000 of the old cavity. At the same time we redesigned the cavity mount to achieve better mechanical stability of the cavity, and thus better length stabilization and quieter transmission measurements during an experiment. With this cavity we implemented a triggered-trapping strategy first proposed some years previously by Scott Parkins and others – namely, the use of strong coupling to trap a single atom with a single-photon-strength field in the cavity [21, 22, 23, 24]. Atomic lifetimes in the cavity were enhanced by a factor of about 4.5 over the free-fall transit time ($\sim 75\,\mu s$ to cross the mode waist), with some rare events lasting longer than 2 milliseconds. During these transits we were able to resolve transmission oscillations associated with atomic motion toward and away from the cavity axis. These transmission signals provided a real-time measurement of atomic position in the cavity. The explicit position information is one-dimensional, but in fact I was able to use it to reconstruct two-dimensional trajectories by exploiting some knowledge of the effective potentials involved. This experiment and the trajectory reconstruction constituted the "atom-cavity microscope," as presented in [25].

At the end of this experiment Christina and I spent a fair amount of time understanding the potentials and heating rates in our system, and how they translated to the qualitative character of trapped atom trajectories and our ability to observe them [26]. In this effort we were aided greatly by Andrew Doherty, and also by Kevin Birnbaum who was a brand-new graduate student at the time. Our immediate thoughts on the experimental front were twofold. One, extend our measurement bandwidth and try to see some signatures of axial motion (along the standing-wave direction of the cavity field, i.e., along the cavity axis). Two, use our real-time position tracking to turn on some feedback and actively cool the motion of each trapped atom during its stay in the cavity. Neither naive attempt worked out very nicely, either in my first attempts with Christina or in those I continued with Kevin after her departure. There were good reasons for this, but before we had a chance to become more educated, the cavity and vacuum system began to exhibit diverse sicknesses.

Once again, then, the experiment was rebuilt, this time more drastically than in the previous cycle. Improvements to the vacuum system, the laser/cavity stabilization scheme, and the data processing system accompanied the replacement of the physics cavity itself. Various technical demons plagued this redesigning effort, such that the experiment is just coming on line at the time of this writing, in April 2003.

Meanwhile, I have investigated the proposed active feedback in experimental conditions and from a more control-theoretic point of view. This work has led to a more carefully reasoned radial feedback algorithm, which has been developed for experimentally realistic conditions and extensively simulated for those same conditions. It has led at the same time to a better understanding of the figure of merit associated with our proposed feedback technique, the limits that apply to it, and directions for future work that should extend our active cooling capability and lead atomic position control to the land of the quantum at long last.

1.5 Overview and Outline

In the preceding sections I have begun with some broad comments on cavity QED as a scientific tool and on the motivations for atomic position control in this setting, as well as a summary of my involvement in research furthering this goal. In this final section of Chapter 1, I will outline the structure and contents of the remainder of this thesis.

In Chapter 2, I develop a basic picture of cavity QED both intuitively and quantitatively. My goal in this discussion is to elucidate the important critical parameters for "quantum" behavior in cavity QED and to show how they appear in the formalism, in historical progression to strong coupling, and in intuitive arguments.

Chapter 3 presents an overview of experimental techniques and ongoing technical issues involved in our single-atom, single-photon experiments. While some experimental details are thoroughly dealt with elsewhere [27, 19], the basic techniques are laid out here in enough detail to define the experimental procedure and address some common points of confusion.

I next move on to results from the atom-cavity microscope experiment in Chapter 4. Again, experimental results and some analysis were provided in [19], but I present a more complete discussion of the reconstruction algorithm for two-dimensional trajectories. The treatment here addresses validation of the algorithm, physical conditions and dynamical regimes that make it possible, and signal-to-noise and sensitivity limits governing the position measurement. Attention is given to benchmarks denoting the degree to which these measurements were truly in a classical regime for the atomic center-of-mass motion.

From the real-time position tracking of Chapter 4 I move on to algorithms for active feedback in Chapter 5. Here I discuss the challenges of active cooling from a control systems point of view, and present simulations building up a viable technique for circularizing atomic motion at a constant distance from the cavity axis. Simulations not only explore highly realistic experimental conditions, but also extend into more hypothetical dynamics in order to clarify the feedback performance and limits more fully.

Chapter 6 describes the feedback experiment currently in progress to implement the strategies derived in Chapter 5. In this chapter the current setup is treated in considerably more detail than the general discussion of Chapter 3. Main new features of this experiment include a new physics cavity, new differentially pumped chamber with a double MOT setup, use of a separate "locking laser" to stabilize the cavity length, and introduction of digital (FPGA) techniques for data acquisition and feedback control.

Chapter 7 is devoted to design and characterization of Fabry-Perot cavities for strong-coupling cavity QED applications. This material is intended both as a technical resource for future design and as a more concrete complement to critical parameter and signal-to-noise considerations developed in Chapters 2 and 4.

Finally, Chapter 8 presents some ideas and calculations for extensions of the active position control work in several directions. I dwell briefly on two or three main avenues for extension: the breaking of cylindrical symmetry to remove atomic angular momentum, the detection and cooling of axial as well as radial motion, and the separation of trapping and sensing to facilitate long lifetimes and effective feedback. The thesis concludes with these discussions of future prospects.