

Copyright
by
Windell Haven Oskay
2001

The Dissertation Committee for Windell Haven Oskay certifies that this
is the approved version of the following dissertation:

ATOM OPTICS EXPERIMENTS IN QUANTUM CHAOS

Committee:

Mark G. Raizen, Supervisor

Michael F. Becker

Roger D. Bengtson

Manfred Fink

Linda E. Reichl

The labour we delight in physics pain.

—Shakespeare, *Macbeth*

ATOM OPTICS EXPERIMENTS IN QUANTUM CHAOS

by

WINDELL HAVEN OSKAY, B.A.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

December 2001

Acknowledgements

It has been my privilege and pleasure to work in Mark Raizen's laboratory. Mark has nurtured a well-funded and smooth-running lab, always full of bright and dedicated people. Mark is also an expert at what is possibly the most difficult problem in physics: choosing interesting problems to study. Beyond this, his neverending creativity and support have been essential in conducting the experiments that are described in this dissertation. More than anything else, I would like to thank Mark for having more confidence in me than I did.

I had the good fortune to work with Daniel Steck on all of the experiments that are described in this dissertation. Dan is in many ways the most talented person that I have ever met, and I have learned a lot from him. Dan performed most of the theoretical and computational work that we did. His skill, efficiency and intuition in these areas is astounding. Dan is equally talented in the laboratory, and I enjoyed working with him on the experiments. Dan has also been a good friend outside the lab, and we have frequently collaborated on experiments with movies, cooking, and chocolate.

Dan and I worked with Bruce Klappauf for most of the kicked-rotor experiments. Bruce was responsible for much of the initial construction of the cesium experiment, and was a wonderful person to work with. Our postdoc Valery Milner also worked on some of the kicked-rotor experiments, and was particularly helpful during some of our longer data runs. More recently, Jay Hanssen has been helpful in some of our longer data runs. The next generation of the experiment will be led by Jay, along with Todd Meyrath and Chuanwei Zhang. The experiment is in capable hands.

Besides the people that worked on the cesium experiment, there have been many other students in our laboratory that I never had the opportunity to work with directly. Kirk

Madison gave me my first tour of the Raizen lab, when I was a prospective student in the spring of 1995. Later, we were housemates, and collaborators on a large number of projects, some of which were work-related. Kirk has an excellent sense of physical intuition, which we have all benefitted from. It seems like forever since Cyrus Bharucha graduated, but he has remained a good friend. I have also enjoyed the company of the other students on the sodium experiment, Martin Fischer, Braulio Gutierrez, Artëm Dudarev, and Kevin Henderson. Our exceptional German students, Artur Widera, Michael Goegler, Nicole Helbig and Alexander Mück have kept the lab colorful and full of chocolate. We have also had a set of dedicated undergraduate students, Wes Campbell, Patrick Bloom, Greg Henry, and Joshua Elliot. Finally, our administrative assistants Adrienne Lipoma and Julie Horn worked very hard to keep our lab running smoothly.

I had three exceptional classes during my time at UT Austin, taught by professors Daniel Heinzen, Thomas Griffy and Allen MacDonald. I would like to thank them for setting such a fine example.

Mike Kash was my undergraduate advisor at Lake Forest College, and it was in his laboratory that I learned how much I enjoy aligning mirrors. I thank him for his continued guidance and encouragement, throughout my time in graduate school.

The various experiments that are described in this dissertation were funded by the R.A. Welch Foundation, the National Science Foundation, the U.S-Israel Binational Science Foundation, and the Sid W. Richardson Foundation. I was supported for one academic year (2000-2001) by a Hutchison Endowment Fellowship, UT Austin.

I would like to thank Roger Bengtson, Michael Becker, Jay Hanssen, Todd Meyrath, Daniel Steck, and Lenore Edman for careful proofreading and corrections in the preparation of this dissertation.

Finally, to my family and friends (and Anderson's Coffee), let me simply say thank you.

ATOM OPTICS EXPERIMENTS IN QUANTUM CHAOS

Publication No. _____

Windell Haven Oskay, Ph.D.
The University of Texas at Austin, 2001

Supervisor: Mark G. Raizen

We experimentally study the quantum dynamics of several simple systems by observing the center-of-mass motion of cold atoms in time-dependent optical potentials. One of the most interesting types of quantum systems is one where the motion in the classical limit is chaotic. Chaotic motion in such a system is inhibited through dynamical localization. Quantum and classical motion in the system exhibit distinctly different overall behavior. One prototypical classically chaotic system is the quantum kicked rotor, which is a system that we can study with our experiments. We have used our setup to observe several phenomena related to dynamical localization, including the dependence of the diffusion rate on short-term correlations. A second topic of study is the effects of noise upon this system. Dynamical localization is a coherent quantum effect that can be destroyed by the introduction of noise, which is an interaction with the external environment. It has been suggested that such coherence-breaking interactions may be necessary for fundamentally quantum systems to exhibit classical dynamics. When we add high enough levels of noise to our experiments, we have found that certain features of the system behavior become indistinguishable from that of a classical system. Beyond the experiments with the kicked rotor, we have studied quantum transport in mixed phase space. We report the first direct observation of chaos-assisted tunneling, where quantum tunneling between two islands in phase space is accelerated by a region of chaotic dynamics that separates them. These experiments required the development of new methods of quantum state preparation, which

we describe here. Finally, we have demonstrated a new method of spatially focusing atoms that has applications in atom lithography.

Table of Contents

Acknowledgements	vi
Abstract	viii
List of Figures	xiii
Chapter 1. Introduction	1
1.1 A Brief History of the Experiment	1
1.2 Classical Motion	3
1.2.1 Regular motion in phase space	3
1.2.2 Chaotic motion: the kicked rotor	6
1.2.3 Momentum diffusion in the kicked rotor	15
1.3 Quantum Chaos	22
1.3.1 The quantum kicked rotor	22
1.3.2 Dynamical localization	24
1.3.3 Quantum resonance	26
1.4 Atom Optics	29
1.4.1 An atom in a standing wave of light	31
1.4.2 Spontaneous scattering	36
1.4.3 Scaled units for the kicked rotor	37
1.4.4 Quantum motion in the pendulum potential	38
1.5 Experimental Context	40
Chapter 2. Experimental apparatus and methods	43
2.1 Overview	43
2.2 The MOT and Optical Molasses	46
2.3 Optical Systems	48
2.3.1 Optical sources	54
2.3.2 Laser stabilization and monitoring	59

2.3.3	Optical lattice calibration	68
2.4	Interaction Chamber	71
2.4.1	Vacuum system	72
2.4.2	Magnetic field control	75
2.5	Imaging	77
2.6	Computer Control	80
2.6.1	Design philosophy	80
2.6.2	Implementation	80
2.7	An Experiment in More Detail	85
2.8	Temperature Measurements	89
 Chapter 3. Experiments with the quantum kicked rotor		92
3.1	Overview	92
3.2	Finite Pulses and Dynamical Localization	92
3.3	Effects of Correlations	99
3.4	Quantum Resonances	103
3.5	Decoherence and the Effects of Noise	110
3.6	The Classical Limit	118
3.6.1	Amplitude noise effects on kicked rotors	118
3.6.2	Systematic effects and the classical model	121
3.6.3	Comparison of the experiment to the classical model	126
3.7	Timing Noise	133
 Chapter 4. Quantum state preparation		140
4.1	Overview	140
4.2	Three Dimensional Lattice Cooling	141
4.3	Stimulated Raman Velocity Selection	149
4.3.1	Overview	149
4.3.2	Experimental setup for velocity selection	153
4.3.3	Optical pumping and pushing	160
4.4	State Preparation by Atomic Sliding Motion	164
4.4.1	Optical phase control	170

Chapter 5. Chaos-assisted tunneling	174
5.1 Overview	174
5.2 Experimental Considerations	178
5.2.1 Amplitude-modulated pendulum system	178
5.2.2 Symmetry requirements	179
5.2.3 CAT Experiments	181
5.3 Data and Results	185
5.3.1 Observation of tunneling	185
5.3.2 Sensitivity to island overlap	186
5.3.3 Sensitivity to velocity class	189
5.3.4 Comparison with pendulum	192
5.3.5 High temporal resolution data	195
5.3.6 Sensitivity to well depth	200
5.3.7 Sensitivity to amplitude noise	206
Chapter 6. Spatial focusing in the kicked rotor system	210
6.1 Overview	210
6.2 The Aperiodic Kicked Rotor	216
6.3 Experimental Methods	223
6.3.1 Near-resonant detection	224
6.3.2 Asymmetric kicking measurements	227
6.3.3 Pendulum oscillation measurements	230
6.3.4 Free expansion measurements	232
6.3.5 Well depth calibration	235
6.3.6 Experimental sequence	236
6.4 Data and Results	237
Bibliography	248
Vita	260

List of Figures

1.1	Phase space portrait of the simple pendulum	5
1.2	Standard map surfaces of section, $K = 0.0$ and $K = 0.1$	9
1.3	Standard map surfaces of section, $K = 1.0$ and $K = 10.0$	11
1.4	Three types of trajectories in the standard map with $K = 1$	12
1.5	Sensitivity to initial conditions: regular trajectories	13
1.6	Sensitivity to initial conditions: chaotic trajectories	14
1.7	The diffusion rate $D(K)$	16
1.8	Locations of the accelerator modes	18
1.9	Detail of an accelerator mode	19
1.10	Growth of momentum in accelerator modes	20
1.11	Illustration of Lévy-flight trajectories	21
1.12	Illustration of dynamical localization	25
1.13	Illustration of quantum resonance	28
2.1	Experimental sequence	44
2.2	Term diagram of the cesium D ₂ line	47
2.3	Characterization of freezing molasses	49
2.4	Photograph of a region of the optical table	50
2.5	Essential optical systems	51
2.6	Photograph of the DBR diode laser	55
2.7	Photograph of the repump laser	56
2.8	Saturated absorption spectrum of the repump laser	61
2.9	Saturated absorption spectrum of the molasses laser	62
2.10	Photograph of the interaction chamber	72
2.11	Illustration of optical access to the interaction chamber	73
2.12	Timing diagram for a kicked-rotor experiment	86
3.1	Calculation of the effective stochasticity parameter	95

3.2	The momentum boundary in phase space	96
3.3	Observation of dynamical localization and the boundary	97
3.4	Quantum diffusion curves at various $\hbar k$	101
3.5	Effects of correlations on momentum distribution evolution	102
3.6	Quantum resonance simulations: two initial conditions	104
3.7	Momentum distributions at quantum resonance: experiment	105
3.8	Momentum distributions at quantum resonance: simulation	106
3.9	Noise and dissipation effects on energy evolution	115
3.10	Spontaneous emission effects on momentum distributions	116
3.11	Classical diffusion rate $D(K)$, with amplitude noise	119
3.12	Profile of a single “kicking” pulse	122
3.13	Illustration of systematic corrections to classical simulations	125
3.14	$E(K)$ with noise: experiment and simulations	127
3.15	$E(t)$ with noise: experiment and simulations; $K = 11.2$	129
3.16	Energy vs noise level: experiment and simulations; $K = 11.2$	130
3.17	Momentum distributions: experiment and simulations; $K = 11.2$	131
3.18	$E(t)$: experiment and simulations; $K = 8.4$	132
3.19	Energy vs noise level: experiment and simulations; $K = 8.4$	133
3.20	Momentum distributions: experiment and simulations; $K = 8.4$	134
3.21	Timing noise effects on energy evolution	135
3.22	Timing noise effects on $E(K)$ curves	136
3.23	Energy growth as a function of timing noise level	137
4.1	Configuration of the 3D far-detuned optical lattice	143
4.2	Optical setup for the 3D far-detuned lattice	146
4.3	Simplified term diagram for stimulated Raman transitions	149
4.4	Optical setup for stimulated Raman transitions	154
4.5	Levels and frequencies for stimulated Raman tagging	156
4.6	Optical setup for the optical pumping and push-away beams	164
4.7	Illustration of the SPASM sequence	165
4.8	Experimental sequence for experiments with state preparation	166
4.9	Timing diagram for experiments with state preparation	168
4.10	Phase control setup for the interaction beam	171

4.11	Phase step response of the interaction beam EOM	172
5.1	Surfaces of section for the driven pendulum; various α	180
5.2	Initial conditions in phase space	182
5.3	Observation of chaos-assisted tunneling oscillations	184
5.4	Momentum distributions for tunneling oscillations	184
5.5	Observation of CAT oscillations with $\hbar = 1.04$	185
5.6	Initial conditions at other locations in phase space	187
5.7	Sensitivity of tunneling signal to spatial location	188
5.8	Sensitivity to changes in momentum ladder	189
5.9	Sensitivity of CAT to width of momentum slices	190
5.10	CAT experiment without Raman velocity selection	191
5.11	Initial conditions in the corresponding pendulum phase space	193
5.12	Comparison of CAT oscillations to pendulum motion	194
5.13	Phase space portraits at different phases of the modulation	196
5.14	Motion of the classical island in momentum space	197
5.15	High temporal resolution tunneling, $\hbar = 2.08$, $\alpha = 7.7$	198
5.16	High temporal resolution tunneling, $\hbar = 2.08$, $\alpha = 11.2$	199
5.17	High temporal resolution tunneling, $\hbar = 1.04$, $\alpha = 10.5$	200
5.18	CAT sensitivity to α , $\hbar = 2.08$	201
5.19	CAT oscillations with $\alpha = 8.0$, $\hbar = 2.08$	202
5.20	CAT oscillations with $\alpha = 9.7$, $\hbar = 2.08$	202
5.21	Tunneling rate as a function of α , $\hbar = 2.08$	203
5.22	CAT sensitivity to α , $\hbar = 1.04$	204
5.23	Tunneling rate as a function of α , $\hbar = 1.04$	205
5.24	Modulated pendulum intensity profiles	206
5.25	Effects of amplitude noise on tunneling, $\hbar = 2.08$	207
5.26	Effects of amplitude noise on tunneling, $\hbar = 1$	208
6.1	Atom deposition with a standing-wave lens	212
6.2	Calculation of atomic focusing after a single kick	218
6.3	Comparison of particle and rotor behavior after one kick	219
6.4	Calculation of atomic focusing after multiple kicks: f_c	220

6.5 Calculation of focusing: distribution evolution 221

6.6 Calculation of focusing: individual distributions 222

6.7 Effects of finite detection pulse duration 228

6.8 Pendulum Rotation: classical Monte Carlo simulations 233

6.9 Pendulum Rotation: classical analytic solution 234

6.10 Momentum contrast time evolution, $K = 0.43$ 238

6.11 Momentum contrast time evolution, $K = 1.06$ 239

6.12 Momentum contrast time evolution, $K = 3.2$ 239

6.13 Pendulum oscillation time evolution, $K = 3.2$ 241

6.14 Pendulum oscillation measurements, $K = 0.43$ 242

6.15 Pendulum oscillation measurements, $K = 1.06$ 243

6.16 Pendulum oscillation measurements, $K = 3.2$ 243

6.17 Pendulum rotation measurements, $K = 1.06$ 244

6.18 Pendulum rotation measurements, $K = 0.43$ 245

Chapter 1

Introduction

1.1 A Brief History of the Experiment

This dissertation describes a series of quantum chaos experiments performed over the course of five years. These experiments were carried out using the tools of atom optics: cold atoms exposed to potentials created with light.

I joined Mark Raizen's group in the fall of 1996. At this time, several pioneering quantum chaos experiments had already been performed in our laboratory using sodium atoms [Robinson95b]. The sodium experiment had observed a range of interesting phenomena including *dynamical localization*, the quantum suppression of chaotic motion [Moore94]. The natural evolution of that experiment was hampered by several technical obstacles associated with the use of sodium atoms. As it turns out, the physical and optical properties of cesium atoms made them more suitable for the next set of quantum chaos experiments [Klappauf99].

In January 1997, I began working on the cesium experiment setup, which was still under construction. This was a new experiment being built up by my co-workers Bruce Klappauf and Daniel Steck. I joined at the stage when the vacuum chamber was being assembled. We first trapped cold cesium atoms in April, but the low number of available cesium atoms forced us to break vacuum and it wasn't until August that we had a fully functional atom trap.

Our first priority on this new experiment was to reproduce the results of the sodium experiment. In particular, we planned to use a realization of the δ -kicked rotor to observe

dynamical localization [Moore95]. After that, we would perturb the system to study the destruction of dynamical localization. The effects of noise and dissipation upon this system are particularly interesting because they give us a clear way to study quantum-classical correspondence.

In the middle of October 1997 we were first able to observe dynamical localization, and by the end of the month, we observed amplitude noise-induced delocalization. We continued to study, in advancing levels of detail, the effects of different types of noise and dissipation upon this system for two years, through the fall of 1999. Ultimately, we were able to drive the quantum kicked-rotor system to its classical limit by adding noise [Milner00; Steck00]. The momentum distributions that we observed in this work were also found to be well described by a modern theory of the shape of the quantum diffusion front [Zhong01].

During this period, we also studied other phenomena in the context of the kicked rotor. For example, we observed the behavior of this system in parameter regimes where the effects of short-term correlations can lead to anomalous diffusion [Klappauf98a]. We also revisited quantum resonances, which had been observed in the earlier sodium experiment [Oskay00].

Beginning in January 2000, we turned our attention from the kicked rotor to chaos-assisted tunneling (CAT). In order to observe CAT, it is necessary to prepare an extremely precise and narrow initial condition at a specific location in phase space. To achieve this, we developed and implemented a new method of quantum state preparation incorporating stimulated-Raman velocity selection. Once this technique was developed, we were able finally to observe CAT in January 2001 [Steck01c]. After performing a detailed study of CAT, and how the tunneling rate varied as a function of various system parameters, we also studied the effects of noise on this system.

Finally, we returned to the kicked rotor in the spring of 2001, but with a different emphasis. Whereas all of our previous work with the kicked rotor system studied momentum transport, we instead worked to observe angular (spatial) focusing of the atoms. A spatially focused array of atoms is potentially useful in lithographic applications. This technique

holds the promise of creating a bright source of these atoms while reducing the number of background (non-focused) atoms compared to other methods.

To describe these experiments properly, it is necessary to cover a certain amount of background material. This introductory chapter will describe some of the classical and quantum systems that we study, and introduce some of our general atom-optic methodology. We must first explore some features of classical transport to understand what is different when we move on to quantum transport. About half of the work described in this dissertation is in the context of experiments with the kicked rotor system. Since the kicked rotor system is such a fine setting in which to introduce key concepts that apply to all of our research, we will begin there. In the chapters that follow, we will describe our experiment in general, how it has been adapted to perform the individual experiments, and what we observed in these experiments.

1.2 Classical Motion

1.2.1 Regular motion in phase space

The experiments that we are concerned with study the one-dimensional center-of-mass (CM) motion of atoms in an optical potential. To an excellent approximation, the external potential is conservative, and the atomic sample is dilute enough that atom-atom interactions are negligible. In this limit, we can model each atom in the experiment as an independent and structureless particle in the potential.

Chaos is most easily defined for classical systems, so let us begin there. Consider the behavior of a single classical particle with N degrees of freedom, vector position \mathbf{x} and momentum \mathbf{p} in a conservative time independent potential. The dynamics are specified by the Hamiltonian $\mathcal{H}(\mathbf{x}, \mathbf{p})$, and Hamilton's equations read

$$\dot{p}_i = -\frac{\partial \mathcal{H}}{\partial x_i}, \quad (1.1)$$

$$\dot{x}_i = \frac{\partial \mathcal{H}}{\partial p_i}, \quad (1.2)$$

where the index i runs from 1 to N . Classical physics is deterministic, so in principle, armed

with Eqs. 1.1, 1.2 and the initial conditions $(\mathbf{x}_0, \mathbf{p}_0)$ at $(t = 0)$, we can calculate the full time evolution of the system.

For certain systems there exist N independent isolating constants of motion (integrals of motion), and such a system said to be *integrable* [Reichl92; Lichtenberg92]. For a one-dimensional ($N = 1$) conservative system, we have the total energy $E = \mathcal{H}(x, p)$ as a constant of the motion, and so such a system is always integrable.

A prototypical integrable system is the simple rigid pendulum. In an appropriate set of units with angle variable x and momentum p , we can take the Hamiltonian to be

$$\mathcal{H}(x, p) = \frac{p^2}{2} - \cos x. \quad (1.3)$$

In order to visualize the dynamics, we will construct a *phase space portrait*. The phase space is a $2N$ -dimensional plot that shows the trajectory $(\mathbf{x}(t), \mathbf{p}(t))$ that a particle with given initial conditions $(\mathbf{x}_0, \mathbf{p}_0)$ traces out. In the case of our 1-D system, each particle traces out a smooth curve in the $x - p$ plane. A pendulum phase space with 18 initial conditions is shown in Fig. 1.1. We will use this example to illustrate some of the generic features of phase space portraits. Our pendulum is periodic in position— we will identify the positions $x = \pi$ and $x = -\pi$ with each other. Since the energy is a constant of motion, curves in the phase space represent trajectories of equal energy, and these paths do not cross (with one exception).

Several types of motion are possible in the rigid pendulum. First, there is the familiar oscillatory pendulum type motion. This type of motion is known as *libration*, and traces out closed curves about the *stable fixed point*, $(p = 0, x = 0)$. The stable, or *elliptic* fixed point represents the pendulum, at rest, hanging directly below the pivot. If a small perturbation is applied, the pendulum will begin oscillating about the fixed point. A second type of fixed point, the *unstable*, or *hyperbolic fixed point*, is found at $(p = 0, x = \pm\pi)$. When a pendulum is at its unstable fixed point, it is inverted: at rest directly *above* the pivot. Obviously, it is unstable because any perturbation will cause it to move quickly away from this position. Trajectories with exactly enough energy to reach an unstable fixed




Figure 1.1: Phase space portrait of the simple pendulum. This plot presents some sample trajectories of particles moving under the influence of the pendulum Hamiltonian, Eq. 1.3. Each curve represents the motion that results from an independent initial condition, and is labeled by its color. Librational motion (a) traces closed curves about the stable fixed point (b). The unstable fixed points (c) have the same energy as the separatrix (d), and represent a stationary, inverted pendulum. Trajectories (e,f) outside the separatrix represent continuous rotation, with $p > 0$ (e) and $p < 0$ (f).

point are said to lie along a *separatrix*. Since there are two possible ways to reach the unstable fixed point (clockwise and counterclockwise motion), the separatrix consists of two curves which intersect there. If a pendulum has *more* than enough energy to reach the unstable fixed point, it will reach the inverted configuration with some nonzero momentum. This type of motion is called *rotation*, because the pendulum will then continue to rotate about its pivot without ever reaching zero momentum. The separatrix then, defines an energy E_s which separates bounded trajectories that circulate about the stable fixed point

($0 < E < E_s$) from those which perform complete rotations ($E > E_s$).

1.2.2 Chaotic motion: the kicked rotor

So far, we have described motion in a time-independent, one-dimensional Hamiltonian system, which is necessarily integrable. We will now begin to examine a Hamiltonian that has explicit, periodic time dependence. This change is equivalent to adding an additional degree of freedom to the system. Systems with higher dimensionality (or time dependence) need not be integrable and we may see the emergence of chaos.

The study of chaos is the study of dynamical systems which are, in a practical sense, unpredictable. Chaos is characterized by extreme sensitivity to initial conditions, in that neighboring trajectories rapidly diverge from each other. This leads to a sort of “practical indeterminism,” where the initial state of the system can never be prepared (or measured) well enough to compute the time evolution of trajectories. Chaotic motion also tends to exhibit ergodicity, and gradually fills up some region of phase space uniformly.

A textbook example of a chaotic system is the δ -kicked rotor. This system has been studied in great detail, and has been one of the primary examples used to explore both classical and quantum chaos. Furthermore, it is admirable in its simplicity; it exhibits a wide range of phenomena, and is experimentally accessible. The kicked rotor is a rigid rotor much like the pendulum, however, it is exposed to gravity only periodically. Between these pulses, or “kicks,” it acts as a free rotor, and we will again assume 2π -periodicity in the angle of the rotor. The pulses are in the form of the Dirac δ function with a given kick strength K , and are spaced apart by unit time. The δ -kicked rotor Hamiltonian is then

$$\mathcal{H}(x, p, t) = \frac{p^2}{2} + K \cos x \sum_{n=-\infty}^{\infty} \delta(t - n). \quad (1.4)$$

Here we have a time-dependent, one-dimensional Hamiltonian system, and therefore energy is no longer conserved. The dynamics of this system are completely governed by its single parameter K , which is known as the stochasticity parameter. For $K = 0$, the Hamiltonian becomes $\mathcal{H}(x, p) = \frac{p^2}{2}$, which describes a decidedly dull integrable system. For

$K > 0$, there is a seemingly limitless variety to the types of motion to be found. We can begin to explore the motion by constructing a discrete mapping for this system.

The application of Hamilton's equations (1.1,1.2) to the kicked rotor Hamiltonian (1.4) yields $\dot{p} = K \sin x \sum_n \delta(t - n)$ and $\dot{x} = p$. Our next step is to integrate these equations over a single kicking period to determine x and p just before each kick. Let t_n be the time at which the n^{th} kick occurs, and let $(t_n - \Delta t)$ be a time just before that kick, where $\Delta t \ll 1$. Integrating \dot{p} over one period, we find

$$\int_{t_n - \Delta t}^{t_{n+1} - \Delta t} \dot{p} dt = \int_{t_n - \Delta t}^{t_{n+1} - \Delta t} K \sin x(t) \sum_{n=-\infty}^{\infty} \delta(t - n) dt \quad (1.5)$$

$$\implies p(t_{n+1} - \Delta t) - p(t_n - \Delta t) = K \sin x(t_n) \quad (1.6)$$

$$\implies p_{n+1} - p_n = K \sin x_n, \quad (1.7)$$

where p_n and x_n are the momentum and position just before n^{th} kick. Eq. (1.7) represents the impulse that the rotor receives from an individual kick, $K \sin x$. Following a similar procedure for \dot{x} ,

$$\int_{t_n - \Delta t}^{t_{n+1} - \Delta t} \dot{x} dt = \int_{t_n - \Delta t}^{t_n} p dt + \int_{t_n}^{t_{n+1} - \Delta t} p dt \quad (1.8)$$

$$\implies x(t_{n+1} - \Delta t) - x(t_n - \Delta t) = p_n t \Big|_{t_n - \Delta t}^{t_n} + p_{n+1} t \Big|_{t_n}^{t_{n+1} - \Delta t} \quad (1.9)$$

$$\implies x_{n+1} - x_n = p_n \Delta t + (1 - \Delta t) p_{n+1}, \quad (1.10)$$

and taking the limit $\Delta t \rightarrow 0$, we finally have

$$x_{n+1} - x_n = p_{n+1}. \quad (1.11)$$

The equations (1.7,1.11) for p and x are usually written in the form

$$\begin{aligned} p_{n+1} &= p_n + K \sin x_n \\ x_{n+1} &= x_n + p_{n+1} \end{aligned} \quad (1.12)$$

and form a discrete mapping that is called the *standard map*, or the *Chirikov-Taylor map*. The standard map contains the same physics as the δ -kicked rotor Hamiltonian (Eq. 1.4), and its simplicity lends itself to intuition and computational study. Moreover, many chaotic systems are locally equivalent to the standard map [Chirikov79]. While the model itself is simple, it describes many forms of complex behavior, of which we will examine a few.

To illustrate some properties of the standard map, we will construct a *Poincaré surface of section*. The Poincaré section is the appropriate analogy of the phase-space portrait for such a time-periodic system, and has many features in common. Instead of plotting the entire trajectory of a particle, we only plot its location in the $x - p$ plane once per period, just before each kick. The standard-map phase space is 2π -periodic with respect to translation in both x and p , and the surfaces of section presented here are plotted modulo 2π . As a first example, we plot the standard map for $K = 0$ in Fig. 1.2. In this figure we have a stroboscopically sampled phase space portrait for the free rotor. As in the earlier phase-space portraits, the curves *do* represent trajectories, however, there is not necessarily one curve per trajectory in the surface of section. The continuous curves that emerge are known as invariant tori and are gradually filled in densely. Generally, periodic sampling of the phase space will reveal rational resonances between the sampling frequency and periodic motion in the dynamical system. In the free rotor case that we are plotting, however, the only time variation is the position $x = pt$. So long as the momentum of one of the initial conditions is not accidentally rational with respect to the sampling frequency, the curves will eventually be filled in completely.

Let us next examine the surface of section for $K = 0.1$, also shown in Fig. 1.2. In this case, we see a mildly perturbed rotor. Many of the invariant curves are only “bent” slightly away from their locations in the integrable limit. These curves are known as Kolmogorov-Arnol’d-Moser (KAM) surfaces or tori, traceable directly to curves of constant momentum in the $K = 0$ case. Beyond these smooth curves that span the entire range in position, complete closed curves are also visible. These closed structures are sometimes referred to as *islands*, and act to confine trajectories. As K increases, chaotic regions begin to appear, and some of the KAM surfaces which serve to confine motion begin to break. The last KAM surface spanning the entire phase space breaks at $K \approx 0.97$. After that, there is no barrier to diffusion in momentum on the scale of a phase-space cell, and we say that this is now a *globally chaotic* system.

Since we are primarily interested in cases which are globally chaotic, we will next

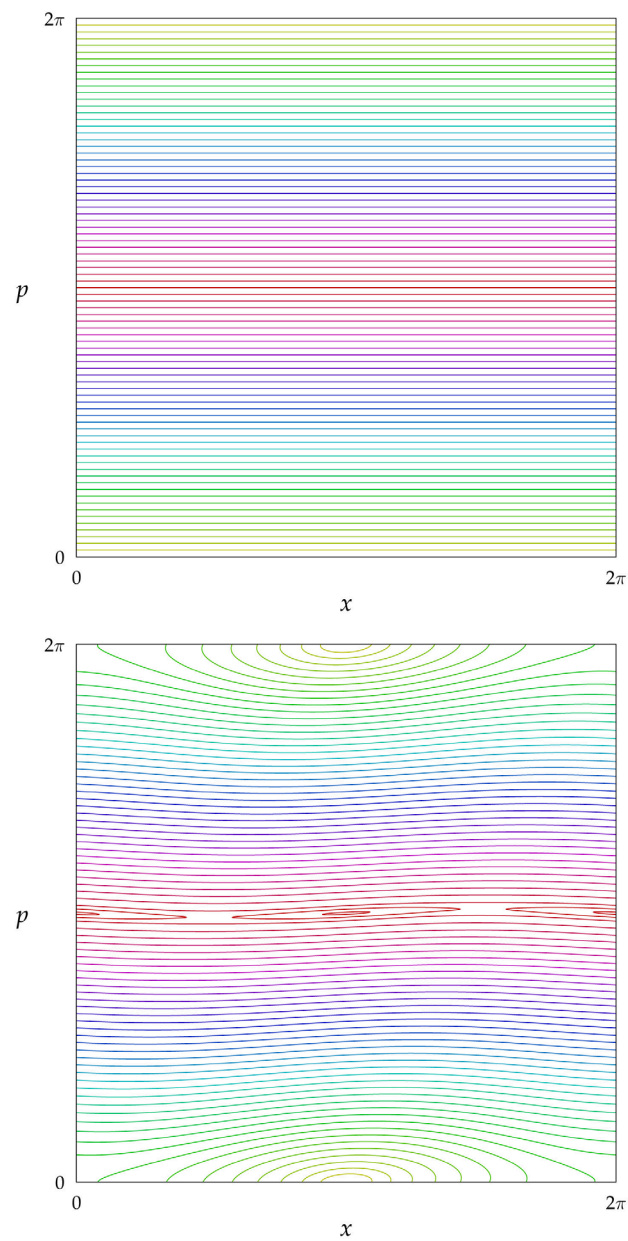


Figure 1.2: Standard map surfaces of section, with $K = 0.0$ (top) and $K = 0.1$ (bottom). Each plot uses about 40 initial conditions; each color corresponds to a different initial conditions. The $K = 0.0$ surface was iterated for about 300000 iterations per point. This is the phase space of an unperturbed rigid rotor. A large number of iterations is necessary to make the KAM tori appear continuous. The $K = 0.1$ surface was iterated for about 8000 iterations per point. The majority of the phase space consists of smooth KAM tori.

examine the surface of section for $K = 1$, shown in Fig. 1.3. This is a mixed phase space, with both regular and chaotic regions. Regions of regular motion are easily identified by the smooth curves that the trajectories trace out. Chaotic motion tends to ergodically (uniformly) fill regions of phase space which are in this case bounded by islands. The apparently random trajectories thus fill these regions up with a “fuzzy” appearance. Again, in the surfaces of section that we present here, each color labels an initial condition. This means that the chaotic regions, through which more than one trajectory typically wanders also take on mixed colors, which do not occur in the regions of regular motion.

As the stochasticity parameter K is further increased, the system becomes increasingly chaotic. Near $K = 2$, we begin to see large, connected chaotic regions that span the entire phase space in both x and p . By $K = 4$, the primary resonances begin to break up, and thereafter the system is predominantly chaotic. For $K > 10$, all of the structure remaining in phase space is too small to be seen without effort. A surface of section for $K = 10$ is plotted in Fig. 1.3.

Our goals in presenting these surfaces of section are to be able to differentiate between regular and chaotic motion, and to develop some intuition for transport properties in the different regions. As a more concrete example, let us examine a few trajectories in the $K = 1$ surface of section in more detail. In Fig. 1.4, we plot clearly separated trajectories in both regular and chaotic regions of phase space. The stable structures that are visible are in the form of regular (KAM) orbits within island chains that consist of pendulum-like periodic resonances. Chaotic regions tend to appear at the (broken) separatrix that joins island chains, and indeed the chaotic region here is visible as a “fuzzy” layer that surrounds an island chain.

How can we be certain that this motion is chaotic? In our brief discussion of chaos thus far, we have noted two chief characteristics of chaotic motion. First, we require that the system exhibits extreme sensitivity to initial conditions. We can quantify this by saying that the distance (in some appropriate measure) between two neighboring trajectories increases exponentially with time. In Figs. 1.5 and 1.6, we compare the sensitivity to initial conditions




Figure 1.3: Standard map surfaces of section, with $K = 1.0$ (top) and $K = 10.0$ (bottom). Each plot uses about 40 initial conditions and about 8000 iterations per point. The $K = 1.0$ plot shows many different types of motion. At this value of the stochasticity parameter, the last KAM surface spanning the phase space is broken. While this phase space is therefore considered to be globally chaotic, large, regular islands are still visible. The phase space for $K = 10.0$ is in contrast nearly homogenous, when viewed at this scale.




Figure 1.4: Three types of trajectories in the standard map with $K = 1$. The standard map (Eq. 1.12) is iterated 1000 times with $K = 1$, for 10 closely-spaced initial conditions in each of three neighborhoods. The three neighborhoods begin at $(x, p) = (3, 3)$, $(1.5, 1.5)$, and $(1.5, \pi)$, and are labeled by color as red, blue, and green, respectively. The ten initial coordinates within each neighborhood are slightly offset in position as $x_i(t = 0) = 1.5 + i \times 10^{-6}$, for $i = 0, 1, \dots, 9$. The trajectories drawn in red and green are examples of regular motion, and fill out invariant tori that are in the form of distinct island chains with different periods. Each individual trajectory stays close to the others in its neighborhood. The trajectories shown in blue exhibit chaotic motion, and fill up a region of the phase space in a seemingly random fashion. Note that these trajectories surround a chain of stable islands. Although each trajectory begins very close to the others, they rapidly diverge. The randomness is not due to the spread in initial conditions: any *single* trajectory in the initial neighborhood also wanders throughout the entire region.

for the regular and chaotic trajectories shown in Fig. 1.4. In the regular case that we have chosen (Fig. 1.5) the distance between neighboring trajectories appears to grow *linearly* as a function of time. In the case that we have identified as chaotic (Fig. 1.6), there is extreme sensitivity to the initial conditions. It is not, however, possible to tell from this figure if we




Figure 1.5: Lack of sensitivity to initial conditions for regular trajectories. The standard map (Eq. 1.12) is iterated 1000 times with $K = 1$, for 10 initial conditions near $(x, p) = (3, 3)$. These are the same ten trajectories plotted in red in Fig. 1.4. The distances $(x_i - x_0)$ between each particle offset in position and the non-offset particle is plotted as a function of the iteration number n , where each curve is labeled by its color. The individual trajectories are distinguishable at all times, and drift apart in a linear fashion. Note that the vertical (position) scale is very small compared to 2π . Since the behavior of neighboring particles is very similar, we say that trajectories in the neighborhood do not exhibit sensitivity to initial conditions.

have exponential sensitivity. We are not presently interested in proving that this system exhibits chaos, but in passing, we outline what is necessary to do so. To show rigorously that a trajectory is unstable requires the calculation of a Lyapunov exponent, which quantifies the rate of exponential separation of neighboring trajectories. A positive exponent shows that the trajectories continue to exponentially separate for all time. The regions of phase space for which all trajectories are unstable is then said to be chaotic.

There is one additional feature that we expect of chaotic systems, which is that the system should display ergodicity over some region of phase space. Although the motion that




Figure 1.6: Sensitivity to initial conditions for chaotic trajectories. The standard map (Eq. 1.12) is iterated 500 times with $K = 1$, for 10 initial conditions near $(x, p) = (1.5, 1.5)$. These are the same ten trajectories plotted in blue in Fig. 1.4. The distances between the trajectories are plotted as in Fig. 1.5. The individual trajectories are not distinguishable (on this scale) for short times, however, the paths eventually diverge wildly. Note that we have plotted the surface of section (from which this graph is derived) modulo 2π . If we did not perform the modulo operation, the differences between particle positions would wander over distances very large compared to 2π over the timescales shown here. Some of the trajectories apparently stay near to each other and display correlated motion. These effects are suggestive of residual structure in the phase space.

we observe for the chaotic trajectory in Fig. 1.4 *appears* to be random, it is clearly bounded. The motion here can be described as diffusion which is halted at several structures in phase space. The chaotic trajectory surrounds, and *does not penetrate*, an island chain. The KAM tori which form the islands are a firm barrier to classical diffusion. Besides this sharp edge, there is also a softer barrier to diffusion in momentum, which is apparent because the trajectory does not span (in p) the entire unit cell of phase space. A probable cause of this is the remaining presence of *cantori* in the phase space, which are the remains of the broken KAM surfaces, and tend to inhibit diffusion to various degrees. For higher values of the

stochasticity parameter K , these types of structures become less important, and diffusion occurs more easily.

1.2.3 Momentum diffusion in the kicked rotor

Many of the properties that we have described thus far are fairly generic, however certain quantitative properties of the kicked rotor are important for our experiments. In the uniformly stochastic phase space of the kicked rotor for large K , we can begin to look at statistical transport properties. A particularly important property is the energy $E_n \equiv \langle p_n^2 \rangle / 2$ imparted to an ensemble of kicked rotors during a sequence of n kicks. Integrating the standard map over n kicks, we find

$$E_n = \frac{1}{2} \sum_{m, m'=0}^{n-1} \langle K \sin x_m K \sin x_{m'} \rangle. \quad (1.13)$$

In the limit $K \rightarrow \infty$, we expect that the motion is nearly random and kick-to-kick correlations can safely be ignored. This assumption is known as the *random phase* approximation, and is also referred to as the *quasilinear* limit. In this case, Eq. 1.13 reduces to

$$E_{n_{\text{ql}}} = \frac{K^2}{2} \sum_{m=0}^{n-1} \langle \sin^2 x_m \rangle = n \frac{K^2}{4}. \quad (1.14)$$

The energy growth is linear in time (number of kicks) in this limit, and we can thus define a quasilinear diffusion constant $D_{\text{ql}} \equiv E_{n_{\text{ql}}}/n = K^2/4$. A secondary signature of diffusive motion, which will become important later, is that the momentum distribution is characterized by a Gaussian probability distribution. In experiments, we typically use the rate of energy growth and the shape of the momentum distribution to characterize the dynamics.

The effects of correlations in the kicked rotor are very important, and it is only in the quasilinear approximation that we can define such a simple diffusion constant. The calculation of a more general diffusion rate $D(K) = E_n/n$ requires consideration of higher order correlations. An analytic expression for $D(K)$ in terms of powers of Bessel functions (to second order) was calculated to be

$$D(K) = \frac{K^2}{2} \left(\frac{1}{2} - J_2(K) - J_1^2(K) + J_2^2(K) + J_3^2(K) \right), \quad (1.15)$$




Figure 1.7: The diffusion rate $D(K)$ is calculated by several methods. First, the smooth curve shown in red is the quasilinear diffusion rate $D_{\text{ql}} = K^2/4$. The green curve is the Rechester and White expression (Eq. 1.15), and oscillates about the quasilinear value. We have also plotted, in blue, the diffusion rate calculated from a simple simulation. A set of $5 \cdot 10^4$ initial conditions were randomly chosen. For each of these, the standard map was iterated 200 times for various values of the stochasticity parameter K . The numerically determined diffusion rate follows the Rechester and White expression very well, except near the maxima, where the simulation shows strong peaks.

by Rechester and White [Rechester80; Rechester81]. This formula is plotted along with $D_{\text{ql}}(K)$ and the results of a simulation in Fig. 1.7, where we see that this expression describes oscillations about the quasilinear diffusion rate.

Eq. 1.15 is intended to model the average, long-time diffusion, and implicitly neglects very-long-time correlations, which at first glance seem unlikely when K is large. However, when comparing the analytical expression to the simulation in Fig. 1.7, it is apparent that there is disagreement at certain values of K . The simulation shows large peaks that are not fully accounted for by the Rechester and White expression. The peaks are due to the

presence of *accelerator modes* in the phase space.

Accelerator modes are small, stable structures in the phase space which tend to rapidly increase the momentum of particles. In the standard map, the largest accelerator modes form stable islands within a range given by

$$2\pi l < K < \sqrt{(2\pi l)^2 + 16}, \quad (1.16)$$

for integer $l \geq 1$ [Chirikov79]. A particle that is trapped inside the accelerator mode hops in momentum by $2\pi l$, monotonically, during each iteration of the standard mapping. The largest set of accelerator modes occurs for $l = 1$, just above 2π . Their location in phase space is shown in Fig. 1.8.

The symmetry of the phase space requires that there be two accelerator modes which accelerate particles in opposite directions. The structure of these islands is an excellent example of the complexity of the standard map, and of one of the islands is shown in detail in Fig. 1.9. The stability of the islands and the nature of the accelerator modes leads to “streaming” behavior, rather than diffusive motion for trajectories in the islands, as illustrated in Fig. 1.10.

In the presence of true accelerator modes, the long-time diffusion rate is expected to diverge because of the streaming behavior. The addition of a small amount of noise to the system breaks the symmetry of the system enough that particles in streaming trajectories eventually fall out of them [Karney83]. The substitution of finite-length pulses for δ -functions in the kicked rotor system has a similar symmetry-breaking effect. In these cases, we say that we have *quasiaccelerator modes* [Lichtenberg92].

Besides the direct streaming behavior in the (quasi)accelerator modes, there is another mechanism by which dramatic changes in momentum occur. The border regions of the accelerator modes are “sticky” in the sense that (otherwise) chaotic trajectories that wander near these regions may become trapped in there and be accelerated to high momentum before wandering away again [Zaslavsky97]. This type of motion is suggestive of *Lévy flights*, which are rare but significant events that strongly influence the overall dynamics




Figure 1.8: This surface of section for the standard map shows the locations of the accelerator modes in phase space for $K = 6.50$. A large number of trajectories generated from random initial conditions are plotted. The color scheme in this surface of section reflects the instantaneous momentum of the particles, rather than the initial condition that it originated from. Particles with high velocity in the $+p$ direction are colored yellow or green, and particles moving quickly in the opposite direction appear as purple or blue. When this color scheme is applied, the accelerator modes appear as brightly colored islands in a dull brown sea of chaos. The islands straddle $p = 0, \text{ mod } 2\pi$. A magnified view of one of these accelerator modes is presented in Fig. 1.9.

[Klafter96]. Examples of this type of motion are shown in Figs. 1.10 and 1.11. The presence of Lévy-flight behavior or quasiaccelerator modes may lead to superdiffusive energy growth. We have also seen that it is possible for energy growth in the system to be limited by the presence of other classical structures. Both of these cases are examples of *anomalous diffusion*, energy growth described by $E(t) \propto t^\mu$, where $\mu \neq 1$ [Chirikov84]. For the relatively short time scales that are used in our experiments, however, the assumption of diffusive energy growth ($\mu = 1$) is appropriate. The diffusion rate expression Eq. 1.15 provides a practical expression for the diffusion rate in the absence of the largest accelerator modes.




Figure 1.9: Detail of an accelerator mode in for $K = 6.50$. This is a magnified region of the phase space displayed in Fig. 1.8, and shows some of the rich detail in these structures. This accelerator mode is an asymmetric island surrounded by a chain of five islands, each of which is surrounded by five more islands. The edges of each main island appear to be bent around the islands that surround it. Some of the particles moving at high velocity (those that have been accelerated, again labeled by their colors) appear in the cloud-like borders of the various islands. The boundaries of the region shown are $0.47 \cdot \pi \leq x \leq 0.69 \cdot \pi$, and $(-\pi/9) \leq p \leq (\pi/9)$.




Figure 1.10: Illustration of momentum growth in the accelerator modes with $K = 6.50$. Three closely-spaced families of initial conditions are used to illustrate the different features of the phase space. The three families begin at $(x, p) = (1.60, 0)$, $(\pi, 0)$, and $(4.4, 0)$, and are labeled by color as green, red, and blue, respectively. Three initial positions within each family are mutually offset by $\Delta x = 0.1$. This places the “green” particles in the positively accelerating mode, the “red” particles in the chaotic sea, and the “blue” particles in the opposite accelerator mode. The momentum of each particle is plotted as a function of time in the standard map. The two groups in the accelerator modes move apart neatly. In contrast, the particles that started in the chaotic region diffuse independently. An interesting feature of this plot is that one of the (otherwise) chaotic trajectories undergoes a long flight, between n of about 650 – 875. The rate of this tremendous acceleration is clearly that of the trajectories shown in blue, and is again suggestive of the “sticky” nature of the accelerator modes. In fact, if we plot this particular trajectory in phase space, it spends a long time in the cloud-like border region (Fig. 1.9) of the negative-going accelerator mode.




Figure 1.11: Illustration of Lévy-flight trajectories. The particle momentum as a function of time is plotted for two trajectories in the standard map with $K = 6.55$, for $5 \cdot 10^5$ iterations. The initial coordinates were chosen as $(x, p) = (0.287, 0.084)$, and $(0.272, 0.084)$. The motion alternates between diffusive, random-walk type behavior and enormous, sudden jumps in momentum.

1.3 Quantum Chaos

Quantum transport is, in general, quite different from classical transport. In the classical picture, we follow the instantaneous momentum and position of particle trajectories through a phase space. Instead, let us now consider the time evolution of a state vector $|\psi\rangle$ under the Schrödinger equation. In quantum mechanics, there is not a well-defined trajectory in phase space, only the evolution of a probability density. What then, if any, is the quantized analogy of a chaotic trajectory? To begin to address this question, we will examine the kicked rotor system in quantum mechanics.

1.3.1 The quantum kicked rotor

The quantum kicked rotor is described by the same physical model as the classical one, however, we now quantize the momentum states. The value of the momentum is constrained to be $j\hbar$, for some integer j . This confines a rotor to a “ladder” of momentum states. This quantization has several possible origins in a physical system. For atoms in a laser field, we may view the ladder as arising from the fact that the atom can only interact with the field in units of whole photons. The quantization may also be taken to arise directly from the spatially periodic nature of the potentials that we employ. An example pendulum potential (which could be pulsed on to realize the kicked-rotor potential) is $V(x) = \alpha \cos(x)$. This potential is spatially periodic, since $V(x + 2\pi) = V(x)$. Bloch’s theorem tells us that the solutions ψ to the Schrödinger equation for this potential satisfy

$$\psi(x + 2\pi) = e^{2\pi ik} \psi(x). \quad (1.17)$$

Furthermore, the topology of our rotor specifies that $x + 2\pi \equiv x$, and therefore $\psi(x + 2\pi) = \psi(x)$. This requirement means that $e^{2\pi ik} = 1$, and so $k \in (0, \pm 1, \pm 2 \dots)$. In physical units where the displacement is a distance (and k is the wave number), this quantization condition specifies the amount of momentum that may be exchanged with a particle. The momentum ladder for an ensemble of particles (as opposed to rotors) is similarly quantized, however each particle may have its own ladder $p = p_0 + j\hbar$, where $0 \leq p_0 < \hbar$ is a constant for each particle.

The quantum kicked rotor Hamiltonian is a straightforward generalization of the classical kicked rotor Hamiltonian (Eq. 1.4) obtained by making the substitution $p \rightarrow \frac{\hbar}{i} \frac{\partial}{\partial x}$. The dynamics are then given by the Schrödinger equation,

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \mathcal{H} |\psi\rangle. \quad (1.18)$$

In order to analyze the behavior of this system, we would like to follow a similar procedure to that which we used in the classical case. We begin by charting the evolution over a single kicking period. This evolution consists of a single kick at time $t = n$ followed by free evolution until the time $n + 1$ at which the next kick would occur. We integrate over the delta kick to find the “kicking” operator $U_0(K) = \exp(-\frac{i}{\hbar} K \cos x)$ and then operate with the usual free-evolution term. These parts may be written together as the single operator

$$U(n + 1, n) = \exp\left(-\frac{i}{\hbar} \frac{p^2}{2}\right) \exp\left(-\frac{i}{\hbar} K \cos x\right), \quad (1.19)$$

so that the state vector at time $t = n + 1$ is given by $|\psi(n + 1)\rangle = U(n + 1, n)|\psi(n)\rangle$. Naturally, repeated application of this operator may be used to find the state vector at later times.

The kicked rotor Hamiltonian is time-periodic and the natural basis set in terms of *Floquet states*. The Floquet theorem is the temporal analog of the Bloch theorem for spatially periodic systems. Since the Hamiltonian is time dependent, the basis states are not energy eigenstates, but rather *quasienergy* states. The basis Floquet states are time-dependent eigenfunctions of the single-kick evolution operator U ,

$$U(n + 1, n)|\psi_{\mathcal{E}}\rangle = \exp\left(-\frac{i}{\hbar} \mathcal{E}\right) |\psi_{\mathcal{E}}\rangle, \quad (1.20)$$

where the quasienergies are identified as the values \mathcal{E} . Since the Floquet states form a complete basis, we may expand an arbitrary initial state $|\psi_0\rangle$ as

$$|\psi_0\rangle = \sum_{\mathcal{E}} A_{\mathcal{E}} |\psi_{\mathcal{E}}\rangle, \quad (1.21)$$

and the subsequent time evolution is limited to states with nonzero $A_{\mathcal{E}}$.

1.3.2 Dynamical localization

The fact that the kicked-rotor system can be described by simple mappings in both classical and quantum physics (Eqs. 1.12 and 1.19) has made it the object of extensive numerical study [Reichl92; Lichtenberg92]. If we compare the energy growth of classical and quantum kicked rotors, we see something very interesting. Recall that the classical energy growth is expected to be diffusive so that $E(t) = Dt$. In the simple simulations shown in Fig. 1.12, we see that this describes the classical behavior quite well. In contrast, the quantum system follows the classical system only for a short time, after which the diffusion is dramatically reduced. This phenomenon is known as *dynamical localization*, and was first identified by Casati, Chirikov, Izrailev, and Ford [Casati79]. It was later shown by Fishman, Grepel, and Prange that the quantum kicked rotor problem can be formulated to resemble the Anderson problem of quantum transport in a disordered lattice [Fishman82; Anderson58; Altland96].

Suppose that we have motion in a perfectly ordered lattice. Then the Bloch theorem gives us solutions (Eq. 1.17) that consist of plane waves modulated by periodic functions. These Bloch states are special solutions for which ballistic motion through the crystal is permitted. On the other hand, if the lattice is not *exactly* periodic, these solutions should not be expected to exist. *Anderson localization* is a phenomenon in condensed matter physics concerning the motion of electrons in a disordered medium. In this case, the electron wave functions become exponentially localized in position, and the material may be viewed as an insulator. In dynamical localization, momentum plays the role of position, and the disorder is introduced by the (pseudorandom) dynamics, rather than by lattice defects. The Floquet states for the kicked rotor are thus exponentially localized in momentum.

We have seen that an arbitrary initial state can be expanded in the Floquet basis (Eq. 1.21). If the initial state is sufficiently narrow, then it must be projected into a coherent superposition of several states. These states evolve independently and eventually dephase. The long-time limit of the momentum distribution then becomes an incoherent sum of the basis states that are populated, and is thus exponentially localized. For short enough times,




Figure 1.12: Illustration of dynamical localization. The energy $\langle p^2/2 \rangle$ is plotted as a function of time for classical (red) and quantum (blue) kicked rotors with $K = 10$. The classical rotor evolution is given by the standard map Eq. 1.12 with 10^3 randomly chosen initial conditions, and the quantum evolution is given by Eq. 1.19, with $\hbar = 1$, for an ensemble of 17 minimum-uncertainty wavepackets uniformly distributed in position.

the discrete nature of the (populated) quasienergy states cannot be resolved (by energy-time uncertainty), and the system energy is expected to grow diffusively. After a characteristic time, the *quantum break time*, the momentum distribution begins to approach its long-time appearance and the energy of the system no longer increases.

It is apparent that the behavior in our classical and quantum kicked rotors is quite different. In the classical system, the momentum distributions are Gaussian, and the system energy grows linearly with time. In the quantum system, the system only evolves diffusively for a short time. After that, the momentum distributions are exponential and the energy does not grow at all. We may interpret the results in the following way: quantum mechanics

suppresses chaotic motion.

Although we have used one specific physical system to demonstrate these results, both chaotic motion and dynamical localization are generic phenomena. Dynamical localization arises from the time-periodic nature of the Hamiltonian, and so we should expect similar behavior in other driven systems as well. This also provides an answer to what we mean by “quantum chaos.” Although there is not chaos in the long-time dynamics of the quantized version of classically chaotic system, this is precisely what we mean. “Quantum chaos” is, for our purposes, the study of classically chaotic dynamical systems.

1.3.3 Quantum resonance

The classical kicked rotor is controlled entirely by the value of the stochasticity parameter K . In the quantum kicked rotor, there are two parameters, K and \hbar . Naturally, we cannot alter one of the fundamental constants of the universe, but we *can* change the relative action scale of certain experiments by wise parameter choice. For the purpose of this discussion, we will assume that the unit of quantum action is determined by a dimensionless, scaled Planck constant \hbar . In our experiments, \hbar is usually determined by the period of the external modulation.

Recall the single-kick time evolution operator, Eq. 1.19, which we have used to describe the quantum kicked rotor. The operator consists of two operations, a kick and a free evolution term described by

$$U = \exp\left(-\frac{i}{\hbar} \frac{p^2}{2}\right). \quad (1.22)$$

Suppose that our initial condition consists of a plane wave with momentum p_0 . During a kick, the particle is only coupled to the other plane waves on its momentum ladder, those with momentum $p = p_0 + j\hbar$, for integer j . If the momentum ladder is defined by $p_0 = 0$, then a particular choice of \hbar can lead to a free evolution term of $U = \pm 1$ for every state on the momentum ladder. This condition is called the *quantum resonance*, and occurs any time that \hbar is an integral multiple of 2π [Izrailev79]. This is similar to the Talbot effect in

optics [Berry99]. If $\hbar k$ is an even multiple of 2π , the free-evolution factor U is unity. If $\hbar k$ is an odd multiple of 2π , the free-evolution factor can either become unity or alternate signs between successive kicks, if the proper momentum ladder is chosen. (It is also possible to choose a momentum ladder that does not display such strong behavior.) In general, if $\hbar k$ is any *rational* multiple of 2π , some sort of quantum resonant behavior is expected. The timescales for the observation of the general quantum resonances is expected to be much longer than for the simplest cases, and we will not discuss them further.

When the free evolution factor U is unity, (e.g., when $\hbar k = 4\pi$), the single-kick and free evolution time evolution operator Eq. 1.19 reduces to

$$U(n+1, n) = \exp\left(-\frac{i}{\hbar} K \cos x\right), \quad (1.23)$$

which we recognize as the operator $U_0(K)$ that describes the kick by itself. For a sequence of N kicks, we apply this operator N times and have

$$U(n+N, n) = \left(\exp\left(-\frac{i}{\hbar} K \cos x\right)\right)^N = \exp\left(-\frac{i}{\hbar} NK \cos x\right), \quad (1.24)$$

which is equivalent to a single-kick operator where the new kick strength is $N \cdot K$. This tells us that because the evolution between kicks does not affect the system ($U = 1$), a long sequence of kicks “collapses” into a single strong kick. The time evolution is simply governed by the fact that the overall kick strength grows linearly with the number of kicks. We can illustrate this behavior by considering the effect of a δ -function pulse on a plane wave with $p_0 = 0$. After a single pulse, the probability of being diffracted into the momentum state $p = j\hbar k$ is

$$P_j = J_j^2(NK/\hbar k), \quad (1.25)$$

where $J_j(x)$ is an ordinary Bessel function [Martin87; Bharucha99]. The Bessel functions $J_j(x)$ peak when x is near j , and thus form a ballistic peak in momentum that moves out linearly with time, as is demonstrated in Fig. 1.13. The system exhibits ballistic growth in momentum where the energy grows quadratically in time. This acceleration is similar in ways to motion in the accelerator modes that occur in the classical standard map. We must




Figure 1.13: Illustration of quantum resonance. These plots show the calculated momentum distribution for the δ -kicked rotor with a plane-wave initial condition and parameters $\hbar k = 4\pi$ and $K = 37$. Two peaks move away from $p_0 = 0$ with nearly uniform acceleration. The distributions shown are calculated from Eq. 1.25. The evolution of the momentum distribution is shown as both a surface plot (left) and as a density plot (right). The intensity scale for these plots is linear.

exercise caution in this analogy because the origins of these two phenomena are entirely different.

A case is also possible where the free-evolution factor U is equal to -1 for a single kick. Then, for a sequence of N kicks, the *total* free-evolution factor $U_{\text{tot}} = U^N = (-1)^N$ alternates sign between successive kicks. This occurs when $\hbar k$ is an odd multiple of 2π , for an appropriately chosen initial momentum state. When this happens, the effect of the second kick is to “undo” the effect of the kick before it. This condition is sometimes called a quantum *antiresonance*. The initial momentum state is critically important if we wish to resolve this type of motion. In our experiments, we never have exactly a plane wave, and we will see later that the continuous nature of our initial conditions can lead to very different behavior. A more detailed theoretical consideration of the dynamics can be used to account for both momentum ladders with $p_0 \neq 0$ as well as the initial momentum distribution of the atomic sample [Bharucha99]. Independent of these considerations, atoms initially near $p = 0$ will undergo ballistic motion as we have described, for $\hbar k = 4\pi$. The overall behavior

leads to momentum distributions characterized by a localized central region surrounded by ballistic peaks, as we will see later.

The evolution of the system near the quantum resonance is quite different from both the classical chaotic motion and the dynamical localization that we saw earlier. Although we did not mention it earlier, it is the case that dynamical localization occurs when quantum resonances do not: when $\hbar k$ is an irrational multiple of 2π . Both dynamical localization and quantum resonances are purely quantum coherent effects, with no analog in *classical* mechanics¹. We have already identified our time-periodic system as being closely related to the system of electron wavefunction propagation in solids. Here, dynamical localization (in momentum) plays the role of Anderson localization (in position). Similarly, the ballistic motion through the momentum lattice at the quantum resonance is much like ballistic motion for Bloch states in a crystal lattice. Since these are coherent effects, it will only be possible to observe them in a system where the fragile quantum coherences are preserved over a sufficiently long timescale.

1.4 Atom Optics

Atom optics is the emerging science of manipulating atomic de Broglie waves. Much of the work in this field has been devoted to developing tools analogous to those that are used with visible light. Mirrors [Balykin88], beam splitters [Moskowitz83], lenses [Friedburg51; Bjorkholm78], and diffraction gratings [Martin87] for matter waves have been demonstrated. The most common methods presently involve the manipulation of neutral atoms by optical fields. There have also been, for example, demonstrations of diffraction through matter gratings [Keith91], and manipulation with magnetic fields [Friedburg51; Migdall85]. An independently evolving set of tools exists for manipulating ions [King99], which are strongly affected by electric fields.

¹There are appropriate analogies in certain wave-mechanical classical systems, such as acoustic or optical systems. As an example, note that the kicked rotor system is much like the propagation of light through a series of diffraction gratings.

We are, in this dissertation, concerned only with the optical manipulation of cold neutral atoms. The many potential applications of this technology include improved atomic clocks [Gibble93], atom interferometry [Kasevich91], a wide range of fundamental physics experiments (such as those performed in our laboratory) [Madison98], and atom deposition on surfaces [Timp92]. The development of laser cooling and trapping in the 1980's led to the creation of now-standard laser cooling techniques such as the magneto-optic trap (MOT) [Raab87]. The MOT uses a combination of near-resonant laser light and magnetic fields to reduce the temperature of an alkali gas sample to very low temperatures—typically $10\ \mu\text{K}$ for cesium atoms. At such low temperatures, the average momentum of the atoms is of the same order of magnitude as the *recoil momentum*, the momentum that an atom acquires by scattering a single photon. (For a cesium atom, the recoil velocity is about $3.5\ \text{mm/s}$.)

Once the typical atomic velocity is this low, the momentum imparted by a single photon scattering off of the atom can become significant compared to the total atomic momentum. It is only at these low momentum values that matter is easily manipulated by optical forces. Any macroscopic object is heavy enough that the momentum imparted to it by a single photon recoil is negligible.

There are, generally speaking, two types of forces that occur when atoms scatter photons. First, there is a dissipative, incoherent scattering force that results from spontaneous emission. Under certain circumstances, an atom may absorb a photon, remain in the excited state for some characteristic lifetime and decay by emitting a photon in a random direction. We will refer to this as the spontaneous force F_{sp} , sometimes called *radiation pressure*. Spontaneous forces are most important very close to atomic resonance, and these are the forces that we use to trap and cool our atomic sample in the MOT. There is also a conservative, coherent scattering force that results from the interaction of an induced dipole moment of the atom with the electric field of light. This *dipole force* is used to create all of the interaction potentials that we use for our experiments. The strength of the dipole force is proportional to the gradient of the intensity of light. The force can be viewed to arise from the shift in the atomic energy levels induced by the electric fields of the light, which is

referred to as the ac Stark shift.

The *optical lattices* that we employ (usually) consist of one-dimensional retro-reflected laser beams. The interference pattern in space is an intensity grating that varies from bright to dark and bright again in half a wavelength of the light (426 nm). The gradient in the electric field strength creates a spatially periodic potential that can serve as a diffraction grating for atoms [Martin87]. The first quantum chaos experiments in this system were proposed ten years ago by Graham, Schlautmann, and Zoller [Graham92]. They proposed sending a beam of atoms through such a standing-wave diffraction grating that was phase modulated in time. Intuitively, such a time-driven nonlinear quantum system should be expected to exhibit dynamical localization. Furthermore, we might recall that the kicked-rotor system is similar to an atom passing through a series of diffraction gratings. Because of this, we can begin to see the origins of our experimental realization of the kicked rotor.

1.4.1 An atom in a standing wave of light

Let us now explore the origins of the dipole force in greater detail. Although this discussion is perhaps unnecessarily technical since it is available elsewhere, it is nonetheless important. Besides understanding the forces well enough to calculate their values, it is also important to know what assumptions and approximations are implicit in the formulas that we use. This derivation was sketched out in the Graham, Schlautmann, and Zoller proposal [Graham92]. The details have iteratively been filled in by a series of dissertations from our laboratory [Robinson95b; Bharucha97; Klappauf98c; Madison98], and are ultimately available in much greater detail than presented here [Steck01b].

Let us begin our description by examining a two-level atom exposed to a one-dimensional standing wave of light. The atom has ground and excited states $|g\rangle$ and $|e\rangle$ separated in energy by $\hbar\omega_0$, where ω_0 is the frequency of the atomic resonance. The light

field is described by

$$\begin{aligned}\mathbf{E}(x, t) &= \hat{\mathbf{z}}E_0 \cos(k_L x) (e^{-i\omega_L t} + e^{i\omega_L t}) \\ &= \hat{\mathbf{z}}(E^+(x, t) + E^-(x, t)) \ ,\end{aligned}\tag{1.26}$$

where E_0 is the amplitude of one of the traveling-wave electric field components. We have separated the field into two counterrotating components, $E^\pm(x, t)$. The wave number and frequency of the light are given by k_L and ω_L , respectively. This describes a standing wave whose time-averaged intensity is given by

$$\langle I(x) \rangle = c\epsilon_0 \langle \mathbf{E} \cdot \mathbf{E} \rangle = c\epsilon_0 E_0^2 (1 + \cos(2k_L x)).\tag{1.27}$$

The total Hamiltonian for the atom in the field is given by

$$\mathcal{H} = \mathcal{H}_A + \mathcal{H}_{AF},\tag{1.28}$$

the sum of the free-atom Hamiltonian \mathcal{H}_A and the atom-field interaction Hamiltonian \mathcal{H}_{AF} . The free-atom Hamiltonian is the sum of the kinetic energy and the internal energy of the atom:

$$\mathcal{H}_A = \frac{p^2}{2m} + \hbar\omega_0 |e\rangle\langle e|.\tag{1.29}$$

The atom-field interaction is described by

$$\mathcal{H}_{AF} = -\mathbf{d} \cdot \mathbf{E},\tag{1.30}$$

where \mathbf{d} is the atomic dipole operator. We have implicitly used the dipole approximation, which is to assume that the electric field is constant over the extent of our wavefunction.

We can further separate the dipole operator into counterrotating components by writing

$$\begin{aligned}\mathbf{d} &= \mathbf{d}^+ + \mathbf{d}^- \\ &= (a + a^\dagger) \langle e | \mathbf{d} | g \rangle \\ &= (|g\rangle\langle e| + |e\rangle\langle g|) \langle e | \mathbf{d} | g \rangle \ ,\end{aligned}\tag{1.31}$$

where we atomic lowering operator is $a = |g\rangle\langle e|$. We evaluate Eq. 1.30 by invoking the *rotating wave approximation* (RWA), under which we average (to zero) terms that oscillate at twice the optical frequency. We then have

$$\begin{aligned}\mathcal{H}_{AF} &= \mathbf{d}^+ \cdot (\hat{\mathbf{z}}E^-) + \mathbf{d}^- \cdot (\hat{\mathbf{z}}E^+) \\ &= \frac{1}{2} \hbar\Omega (ae^{i\omega_L t} + a^\dagger e^{-i\omega_L t}) \cos(k_L x) \ ,\end{aligned}\tag{1.32}$$

where the maximum Rabi frequency is defined as

$$\Omega = -2\frac{E_0}{\hbar}\langle e|d_z|g\rangle, \quad (1.33)$$

and d_z is the $\hat{\mathbf{z}}$ -component of the dipole operator.

We may, without loss of generality, separate the atomic internal and external degrees of freedom. It is also convenient to transform to a frame that rotates at the laser frequency ω_L to simplify the time dependence of the system. The atomic state vector may be written as

$$|\psi\rangle = |\psi_g\rangle|g\rangle + |\psi_e\rangle|\tilde{e}\rangle, \quad (1.34)$$

where $|\psi_e\rangle$ and $|\psi_g\rangle$ are center-of-mass state vectors. We have implicitly defined a slowly varying excited state

$$|\tilde{e}\rangle \equiv e^{i\omega_L t}|e\rangle, \quad (1.35)$$

which further necessitates defining $\tilde{E}^\pm \equiv e^{\pm i\omega_L t}E^\pm$, $\tilde{a} \equiv |g\rangle\langle\tilde{e}|$, and so forth. Under these transformations, the atom-field Hamiltonian (Eq. 1.32) becomes

$$\tilde{\mathcal{H}}_{\text{AF}} = \frac{1}{2}\hbar\Omega(\tilde{a} + \tilde{a}^\dagger)\cos(k_L x), \quad (1.36)$$

and the free-atom Hamiltonian (Eq. 1.29) becomes

$$\tilde{\mathcal{H}}_{\text{A}} = \frac{p^2}{2m} - \hbar\Delta_L|\tilde{e}\rangle\langle\tilde{e}|. \quad (1.37)$$

The new *effective* excited state energy is determined by $\Delta_L = \omega_L - \omega_0$, the laser detuning from resonance. After making these substitutions, we are now ready to insert $|\psi\rangle$ from Eq. 1.34 into the Schrödinger equation, Eq. 1.18, with Hamiltonian $\mathcal{H} = \tilde{\mathcal{H}}_{\text{A}} + \tilde{\mathcal{H}}_{\text{AF}}$. This substitution yields

$$i\hbar\frac{\partial}{\partial t}(\psi_g|g\rangle + \psi_e|\tilde{e}\rangle) = \left(\frac{p^2}{2m} - \hbar\Delta_L|\tilde{e}\rangle\langle\tilde{e}| + \frac{1}{2}\hbar\Omega(\tilde{a} + \tilde{a}^\dagger)\cos(k_L x) \right)(\psi_g|g\rangle + \psi_e|\tilde{e}\rangle). \quad (1.38)$$

By operating on the left with the *internal state* vectors $\langle g|$ and $\langle\tilde{e}|$ successively, this can be separated into a coupled pair of equations of motion. The equation for the ground state is

$$i\hbar\frac{\partial}{\partial t}\psi_g = \frac{p^2}{2m}\psi_g + \frac{1}{2}\hbar\Omega\cos(k_L x)\psi_e, \quad (1.39)$$

and the excited state equation of motion is

$$i\hbar\frac{\partial}{\partial t}\psi_e = \frac{p^2}{2m}\psi_e + \frac{1}{2}\hbar\Omega\cos(k_Lx)\psi_g - \hbar\Delta_L\psi_e. \quad (1.40)$$

Usually, our experiments take place in a parameter regime where the laser is in the general neighborhood of an atomic resonance, which is necessary for the dipole force to be strong enough. We also require that the laser is detuned far enough that spontaneous scattering is unlikely over the timescales of our experiments. We are after the observation of coherent quantum effects, and so it is critically important that we avoid any incoherent interactions such as spontaneous scattering. In the limit of large detuning where atomic excitation is highly unlikely, we can simplify the equations of motion by eliminating the excited state amplitude through an adiabatic approximation. In order to understand the dynamics, let us compare the time and energy scales that the terms of Eq. 1.40 represent. The kinetic energy terms correspond to several recoil energies. The *recoil energy* $E_r = \hbar^2k_L^2/2m$ is the kinetic energy of an atom with one recoil ($\hbar k_L$) of momentum. The corresponding *recoil frequency*, $\omega_r = E_r/\hbar = \hbar k_L^2/2m$ is about $2\pi \cdot 2.07$ kHz for the cesium transition that we use. The detuned laser field induces oscillation within the atom at the generalized Rabi frequency, $\Omega_g = \sqrt{\Omega^2 + \Delta_L^2}$. In our experiments, the Rabi frequency Ω ranges from zero to the order of 500 MHz, and the laser detuning Δ_L is 6-40 GHz, depending on the particular experiment so that the generalized Rabi frequency is given approximately by $\Omega_g \approx \Delta_L$. An additional important timescale is the natural decay rate of the excited state, which for cesium is $\Gamma = 2\pi \cdot 5.2$ MHz. We are attempting to solve for the center-of-mass motion, which occurs on much slower timescales than those of the internal motion. In the adiabatic approximation, we then assume that the internal motion damps instantaneously by taking $\frac{\partial}{\partial t}\psi_e$ to be zero. We can then solve Eq. 1.40 for ψ_e to simplify Eq. 1.39, the equation for the ground state motion. We then have an expression for the ground state amplitude,

$$i\hbar\frac{\partial}{\partial t}\psi_g = \frac{p^2}{2m}\psi_g + \frac{1}{2}\hbar\Omega\cos(k_Lx)\left(\frac{1}{2}\hbar\Omega\cos(k_Lx)\psi_g/(\hbar\Delta_L - \frac{p^2}{2m})\right). \quad (1.41)$$

If we recall that $\hbar\Delta_L \gg \frac{p^2}{2m}$, we can further simplify this by to

$$i\hbar\frac{\partial}{\partial t}\psi_g = \frac{p^2}{2m}\psi_g + V_0\cos(2k_Lx)\psi_g, \quad (1.42)$$

where we have redefined the zero of potential energy and defined the *well depth* V_0 as

$$V_0 \equiv \frac{\hbar\Omega^2}{8\Delta_L} = \frac{d_z^2 E_0^2}{2\hbar\Delta_L}. \quad (1.43)$$

The dynamics described by Eq. 1.42 are that of the pendulum Hamiltonian (Eq. 1.3) substituted into the Schrödinger equation. We can interpret this to mean that our atom in the standing wave (in an appropriate set of limits) acts as a ground-state, point particle. The well depth is the magnitude of the ac Stark shift that creates the dipole forces. From Eq. 1.43, we note that the strength of the dipole force is proportional to the *gradient* of the intensity ($I \sim E^2 \sim \Omega^2$) and inversely proportional to the detuning from resonance, Δ_L .

An important difference from the *classical* pendulum is that the momentum scale here is quantized because the dipole force arises from coherent scattering. An atom that absorbs a photon from the left may scatter that photon to the right, in a process where it gains no net momentum. It may also scatter the photon back to the left, picking up two photon recoils ($2\hbar k_L$) of momentum in the process. In this picture we again see that the atom can only exchange momentum in multiples of $2\hbar k_L$.

Cesium is far from being a two-level atom. Our lasers are tuned near the cesium D_2 ($6S_{1/2} \rightarrow 6P_{3/2}$) transition at 852 nm, which defines a manifold of ground and excited states for our purposes. There are two ground hyperfine states ($F_g = 3, F_g = 4$), each of which has multiple magnetic sublevels m_F . The four excited state hyperfine levels are also multiply degenerate. (We show a term diagram for the relevant atomic structure in Fig. 2.2.) We can simplify the relevant structure somewhat by beginning our experiment with all of our atoms pumped into a single ground-state hyperfine level and continuing to consider the case of linearly polarized light. If the detuning Δ_L is much larger than the width of the excited state manifold (~ 0.5 GHz), the excited states are effectively degenerate. The total dipole moment summed over the possible excited states becomes independent of the magnetic sublevel, with value $d = 2.2 \cdot 10^{-29}$ C · m. A more detailed consideration that does not assume infinite detuning also shows that a two-level approximation is valid, but the effective dipole moment turns out to depend upon the value of m_F .

1.4.2 Spontaneous scattering

We would like to work in a regime where spontaneous emission can be made negligible. One reason for this is that we are attempting to perform “clean” experiments that involve only conservative forces. The spontaneous force F_{sp} is not conservative— it tends to heat our atomic system and break quantum coherences. We need to preserve quantum coherences for relatively long times (up to 4 ms) in order to observe quantum effects. A second reason is that a spontaneous scattering event can change the internal state of the atom. A cesium atom that begins in the $F = 4$ hyperfine ground state can, after a spontaneous event, end up in the $F = 3$ ground state. Since these two ground states are separated in energy by $\sim \hbar \cdot 9.2$ GHz, this can significantly alter the effective laser detuning Δ_L . A simplistic estimate of the scattering rate R_{sc} can be derived from the excited state equation of motion, Eq. 1.40. In the adiabatic approximation we set $\frac{\partial}{\partial t}\psi_e$ to be zero. This yields a relation between the ground state and excited state populations, $|\psi_g|^2$ and $|\psi_e|^2$. Then, if the ground state population is close to unity, we can estimate the total scattering rate to be the product of the spontaneous decay rate Γ and the excited state population,

$$R_{\text{sc}} = \Gamma |\psi_e|^2 \approx \Gamma \left(\frac{1}{2} \frac{\Omega}{\Delta_L} \cos(k_L x) \right)^2 \rightarrow \Gamma \frac{1}{8} \frac{\Omega^2}{\Delta_L^2}, \quad (1.44)$$

after averaging over the standing wave in position. A more detailed calculation using the optical Bloch equations [Steck01a] yields a steady-state expression

$$R_{\text{sc}} = \frac{\Gamma}{2} \frac{I/I_{\text{sat}}}{1 + 4(\Delta_L/\Gamma)^2 + (I/I_{\text{sat}})}, \quad (1.45)$$

in terms of the *saturation intensity* I_{sat} , which is defined by $I/I_{\text{sat}} = 2(\Omega/\Gamma)^2$. We are most concerned with residual spontaneous scattering that occurs when we are in a far-detuned regime. We then note that the total spontaneous emission rate is directly proportional to the intensity and inversely proportional to the *square* of the laser detuning. Recall that the dipole-force strength only decays as $1/\Delta_L$. Because of this, we can usually tune the laser far enough from resonance to make spontaneous emission negligible while maintaining enough dipole force strength to do our experiments. Since cesium is not truly a two-level

atom, precise calculation of the scattering rate requires consideration of the dipole moment for each possible transition.

1.4.3 Scaled units for the kicked rotor

In the everyday course of laboratory work, we use modern SI units [Taylor95]. In contrast, we have used a system of scaled dimensionless units in sections 1.2 and 1.3. This is often useful in elucidating complex physics without dragging along more constants than is absolutely necessary. Scaled units provide a language for the communication of physics that is independent of, for example, atomic data. Frequently, the scaled units also are the *natural* units of the system. Here, we will work through the choice of units that leads from our result for atomic motion in a standing wave to our realization of the quantum kicked rotor.

Let us begin with Eq. 1.42, which is the Schrödinger equation for a system described by the Hamiltonian

$$\mathcal{H}(x, p) = \frac{p^2}{2m} + V_0 \cos(2k_L x). \quad (1.46)$$

In the kicked-rotor experiments, we modulate the amplitude of this potential by pulsing on a series of short pulses spaced apart by time T , so that we have

$$\mathcal{H}(x, p, t) = \frac{p^2}{2m} + V_0 \cos(2k_L x) \sum_n F(t - nT). \quad (1.47)$$

The pulse function $F(t)$ describes a short pulse of unit amplitude and duration t_p , where $t_p \ll T$. We now define a set of transformations,

$$\begin{aligned} x' &= 2k_L x, \\ t' &= t/T, \\ p' &= p\hbar/2\hbar k_L, \\ f(t') &= F(t)/\eta, \\ K &= (\hbar k_L/\hbar)\eta T V_0, \\ \mathcal{H}' &= (\hbar k_L/\hbar) T \mathcal{H}, \end{aligned} \quad (1.48)$$

where we have defined a pulse integral $\eta = \frac{1}{T} \int_{-\infty}^{\infty} F(t) dt$, such that $\eta \propto t_p$ and $\int_{-\infty}^{\infty} f(t) dt = 1$. Furthermore, we define the constant $\hbar k = 8\omega_r T$, where the recoil frequency $\omega_r = \hbar k_L^2/2m$

contains all of the information about the atom. If we examine the commutation relation

$$[x', p'] = \frac{\hbar k}{\hbar} [x, p] = i\hbar k, \quad (1.49)$$

we find that $\hbar k$ indeed plays the role of a scaled Planck constant that tells us the relative action scale of the system compared to \hbar . Note also that we have now defined the kicked rotor stochasticity parameter K in terms of physical parameters.

We have described the scaled units as being, among other things, the natural set of units for the given physical system. Our time unit is the most basic interval in the system, which is the time between subsequent kicks. The position scale is also natural in the sense that x' is 2π -periodic. We saw earlier that the momentum of a given particle is constrained to a ladder of spacing $2\hbar k_L$. In the scaled units, this is taken into account, and momentum is now exchanged in units of $\hbar k$.

We may use the unit transformations (1.48) to rewrite the Hamiltonian in Eq. 1.47 as

$$\mathcal{H}(x, p, t) = \frac{p^2}{2} + K \cos x \sum_n f(t - n), \quad (1.50)$$

after dropping the primes. We can now take the limit of δ -function pulses where $t_p \rightarrow 0$ and K is constant. In this case we recover the δ -kicked rotor Hamiltonian (Eq. 1.4) for a finite number of pulses. We cannot generate an infinite train of pulses in the laboratory, but we can generate a pulse train that is long enough to explore many regimes of dynamics. A more troublesome problem is that we cannot actually generate δ -function pulses. The δ -kicked rotor has been an object of intense study partly because of the computational ease that results from the δ -function in the equations of motion. As we will see later, the short-but-nonzero length of our pulses leads to a slight dependence of K on the momentum.

1.4.4 Quantum motion in the pendulum potential

In our discussions so far, we have described the classical pendulum potential and how it is modulated to realize the classical kicked rotor. We have explored both the quantum and classical dynamics of the kicked rotor system. We have also described the method that we

use in our experiment to realize a quantum pendulum potential, which is in turn modulated to experimentally realize the kicked rotor. Let us take now a step back and explore a few aspects of dynamics in the quantum pendulum potential.

The Hamiltonian for the quantum pendulum in our experiment is given in physical units by Eq. 1.46. As a first step towards analysis, let us define a set of scaled units for the pendulum system. The scaling is very close to that of the kicked rotor system. The most significant difference is that in the absence of periodic modulation, there is no longer a natural time scale for the system. In this case, we may take $\bar{k} = 1$ for convenience, and thus define the time unit T_u . The unit scalings for the pendulum are defined with

$$\begin{aligned}
 \bar{k} &= 8\omega_r T_u = 1, \\
 t' &= t/T_u = 8\omega_r t, \\
 x' &= 2k_L x, \\
 p' &= p\bar{k}/2\hbar k_L = p/2\hbar k_L, \\
 \alpha_p &= V_0/8\omega_r \hbar, \\
 \mathcal{H}' &= (\bar{k}/\hbar)T_u \mathcal{H} = \mathcal{H}/8\omega_r \hbar,
 \end{aligned}
 \tag{1.51}$$

so that the physical Hamiltonian (Eq. 1.46) may be written as

$$\mathcal{H}(x, p) = \frac{p^2}{2} + \alpha_p \cos x,
 \tag{1.52}$$

after dropping the primes. The scaled well depth α_p is the sole parameter in this system. Note that for some of the discussions relevant to this system, we will continue to use physical units.

An important difference between the quantum and classical dynamics of this system is that the quantum dynamics can exhibit tunneling phenomena, as we will discuss later (in chapter 5). One particular flavor of tunneling, *Bragg scattering*, occurs in the pendulum potential and has been observed in several atom optics experiments [Martin88; Giltner95; Kozuma99]. While we do not directly observe Bragg scattering in our experiments, it is important for us to understand the rate of this process.

Generally speaking, Bragg scattering is a multi-photon, nonclassical process that couples states in certain symmetric momentum classes. The most familiar example involves

scattering off of a condensed matter crystal with certain incident angles. The process arises from the periodic nature of the potential, which we have seen also gives rise to the momentum quantization. In our case, we consider atomic scattering off of a stationary 1D optical lattice. In this case, atoms in an initial momentum state $p = n \cdot \hbar k_L$ (for integer n) are coupled by an integer number of two-photon transitions to the state with opposite momentum. We again consider motion in a far-detuned regime where the atoms remain in the ground state. The first-order Bragg process couples ground-state atoms in the two momentum states $p = \pm \hbar k_L$ with what can be described as a two-photon stimulated Raman transition between motional states. (We will return to discuss stimulated Raman transitions in chapter in chapter 4.) The Rabi frequency for the first order process is given by $\Omega_R = \Omega^2/8\Delta_L$, where Ω is the maximum Rabi frequency that we described earlier in the context of the atomic internal dynamics.

In the case of higher-order Bragg scattering, we consider the $2n$ photon transition between momentum states at $p = \pm n \cdot \hbar k_L$. In our far-detuned regime, the excited and intermediate states are not populated, and this complicated process may be described as a two-level system. The expression for the n^{th} order Bragg scattering rate has been shown to be [Giltner95]

$$\Omega_{B_n} = \frac{\Omega_R^n}{(8\omega_r)^{n-1}[(n-1)!]^2} \ , \quad (1.53)$$

by considering the detunings of the intermediate motional states. Often, this oscillation frequency is very small for high-order Bragg scattering. In our scaled units for the pendulum, the frequency is

$$\Omega'_{B_n} = \frac{\alpha_P^n}{[(n-1)!]^2} \ . \quad (1.54)$$

1.5 Experimental Context

The majority of research in quantum chaos is theoretical, however several physical settings have emerged in which to explore the quantum dynamics of nonintegrable systems. Cold atoms in optical lattices are a relatively new entry to this field, and possess some features

which make them advantageous over other systems. This is certainly not a complete list, and it does not cover many of the references directly relevant to our specific experiments, e.g., noise-induced delocalization. We will discuss these results, later, in the context of the individual experiments.

First, mesoscopic structures created in the context of condensed matter physics have been studied in quantum chaos experiments. These systems exhibit a wide range of phenomena, including for example weak localization and scarring [Marcus92; Wilkinson96]. These experiments are necessarily difficult, because of the short relaxation times that occur in any solid system. The phase coherences that are necessary to observe quantum effects are strongly inhibited by electron-electron interactions and lattice vibration.

Another widely studied system is that of multiphoton Rydberg atom ionization, where dynamical localization has been found to suppress chaotic motion that leads to ionization [Bayfield74; Koch95]. The spectroscopic study of atoms in strong external fields has also been a fruitful system [Main91; Delande01]. These experiments address chaotic motion of the *internal* degrees of freedom— the classically chaotic limit of the electron dynamics. This is very different from the atom optics experiments, which ultimately treat the atom as a point particle. The use of atomic systems (Rydberg or atom-optic) for quantum chaos experiments has a strong advantage. The system can be made to be very isolated. Dilute samples of atoms in a high-vacuum environment can be fairly well isolated from the outside world. A possible criticism of experiments with Rydberg electrons is that the quantum state of the atom can be very difficult or impossible to measure; the only quantity that can be easily measured is the ionization threshold [Delande01].

A third class of experiments investigates wave-mechanical phenomena that are analogous to quantum-mechanical systems. Wave equations similar to the Schrödinger equation exhibit similar phenomena in their solutions, and we might well expect “wave chaos” to be similar in meaning to “quantum chaos.” Two-dimensional microwave cavities in the shape of billiards have demonstrated the electromagnetic analogies of many quantum-chaotic phenomena. As an example, dynamical localization [Sirko00] and signatures of chaos-assisted

tunneling [Dembowski00] have been observed in this system. We have noted previously that the kicked rotor system is similar to light passing through a series of diffraction gratings. Dynamical localization has recently been observed in such a system [Fischer00a].

Our atom-optical system is one of the most promising for the continued study of both quantum chaos and other transport properties. Let us review some of the key strengths of this technology. First, we have the environmental isolation that results from using free atoms in a dilute gas. Dissipative interactions can usually be made minor enough to permit detailed study of quantum dynamics. Secondly, we have a wide degree of experimental control. Our potentials are created dynamically using light gated through acousto-optic modulators. We have a fine degree of control over the amplitude and (to a lesser degree) the phase and frequency of the light. Many different potentials can be created in a 1-D sinusoidal system simply by the choice of modulation. We also have a high degree of control over the initial conditions for our experiments. The internal atomic state can be selected by a variety of optical methods, and the overall atomic sample can typically be prepared at a temperature of less than $1 \mu\text{K}$. With the state preparation techniques that we have developed for the chaos-assisted tunneling experiments, we can center cold samples at arbitrary locations in phase space. This is another advantage over other systems, where it may be difficult or impossible to prepare electronic states for experiments. Finally, the momentum distribution in our system can be directly measured with a ballistic (time of flight) measurement technique.

Chapter 2

Experimental apparatus and methods

2.1 Overview

In this chapter we introduce our experimental setup and methods. We begin by describing a typical experiment and gradually describe the individual components that make up the experimental apparatus. The setup has evolved substantially during the course of the work described in this dissertation. Many of the newer cooling and state-preparation techniques have only been in place for the last year. The apparatus and methods that are relevant to these procedures will be described later, in the chapter on quantum state preparation. Here, we will describe only those components of the setup that apply to all of the experiments. Many of the components of both the newer and older parts of the setup have already been described in detail in the dissertations of my co-workers Bruce Klappauf and Daniel Steck [Klappauf98c; Steck01b]. We will in this chapter review the essential features of the experiment, as well as certain parts that have not been described in detail elsewhere.

A simplified diagram that illustrates our experimental sequence is shown in Fig. 2.1. We begin our experiment by collecting a modest sample of 10^6 cesium atoms in our MOT. The atoms are collected in ultra-high vacuum from the low-velocity tail of dilute cesium gas at room temperature. Typically, our MOT loading time is of order 5 s. After loading the MOT, we apply additional cooling and/or state preparation sequences. For the simplest set of experiments, this consists of detuning the laser light used in the MOT (the *optical molasses*) further from resonance for several milliseconds. In later experiments, we apply a detailed state preparation sequence at this stage.




Figure 2.1: Schematic representation of a typical experimental sequence. We begin by collecting atoms in a standard six-beam magneto-optic trap (MOT). We then turn off this trap and expose the atoms to the one-dimensional interaction potential. A ballistic expansion interval is used to transform the momentum distribution into a position distribution. We measure this distribution by “freezing” the atoms in place in optical molasses and imaging them with a CCD camera. A similar diagram for the more complex experiments involving quantum state preparation is shown in Fig. 4.8

After cooling and state preparation is complete, we turn off the trapping fields and begin the experiment itself. Our one-dimensional optical lattice that creates our interaction

potential is switched on and intensity modulated by an acousto-optic modulator. The total interaction time in most of our experiments is less than 2 ms. During the interaction period, atoms exchange momentum with the standing wave. The initial distribution provided by the MOT is well localized in position and the final momentum distribution can be measured by a time of flight technique. We allow the atoms to ballistically expand in the absence of external fields for 15 ms. During this free drift period, the momentum distribution converts itself into a position distribution, since the fastest atoms move the farthest from the initial MOT position. We then “freeze” the atoms in place with optical molasses. During this freezing period, the atoms scatter light from the optical molasses beams which we image on a charge-coupled device (CCD) camera. We typically integrate the picture from the camera along one axis to produce the one-dimensional momentum distribution.

Our measurement technique is destructive, so each experiment begins with a fresh set of cold atoms in the MOT. For experiments in which we sweep the total interaction time as one of the system parameters, each time step requires an independent experiment. Furthermore, we commonly repeat an experiment many times and average the results together. Repeated measurements are warranted either when the signal-to-noise ratio is very low (for example in the chaos-assisted tunneling experiments), or when we wish to average over a system parameter. In experiments where we added noise to the interaction, we averaged over many independent realizations of the noise.

The execution of this somewhat complex sequence of events requires the cooperation of many experimental systems. Several independent lasers are necessary to provide the light for the atom trap and the far-detuned interaction beam. A large collection of optics and support electronics are used to keep the lasers at the correct frequencies. The interaction itself takes place inside a vacuum system, into which all of the laser beams must eventually be directed. The entire experiment is controlled by a computer that automates the timing and data collection.

2.2 The MOT and Optical Molasses

The magneto-optic trap is in some sense the heart of our experiment: it is the basic tool that we use to gather atoms and cool them. This trap has become a common tool in many branches of modern atomic physics, and it is now described in textbooks [Demtröder96; Metcalf99]. As excellent references are available elsewhere, we will not review it in detail here. The MOT works by using a combination of magnetic fields and laser light. The light is circularly polarized in a $\sigma^+ - \sigma^-$ configuration, typically detuned (at least) several natural linewidths to the red of resonance, and is ideally near the saturation intensity. The particular resonance chosen is typically a *cycling transition*, one where the excited state is most likely to decay back to the ground state that it was excited from. Efficient polarization-gradient cooling [Dalibard89] in this configuration typically leads to ultimate temperatures several times the *recoil temperature*. The recoil temperature is where the average atomic velocity is one recoil velocity. The quadrupolar magnetic fields are generated by a pair of anti-Helmholtz coils. These coils define a magnetic-field zero, near which the atoms in the MOT are confined.

Our implementation of the MOT is the standard six-beam geometry, where three of the beams are formed by retroreflecting the other three orthogonal beams. The beams intersect the magnetic field zero in the center of our interaction chamber (vacuum chamber). All of our lasers are close to the cesium D₂ ($6S_{1/2} \rightarrow 6P_{3/2}$) transition at 852 nm. The specific cycling transition that we use is the $6S_{1/2}, F_g = 4 \rightarrow 6P_{3/2}, F_e = 5$ transition, where F_g and F_e represent the ground and excited state hyperfine sublevels, respectively. The cycling transition is not perfect however, and there is a 0.01% chance per spontaneous scattering event of making an (undesired) transition to the $6S_{1/2}, F_g = 3$ ground state instead. Once in this state, an atom no longer “sees” the cooling light, and the cooling process ceases. Although this may seem like a slow decay, the great number of scattering events in our trap leads to a measured trap lifetime of about 3 ms if this loss process is allowed to occur. A standard method for eliminating this loss is to provide a way of returning atoms to the cycling transition. In our setup, an independent laser, the *repump*




Figure 2.2: Term diagram of the D_2 line of the ^{133}Cs atom. The ground-state hyperfine splitting given here is exact because this so-called “clock transition” between the $6S_{1/2}$, $F_g = 3, 4$ is used as the present definition of the SI second. The excited-state splittings are taken from [Tanner88]. (Diagram: D. Steck.)

laser, is tuned directly to the $6S_{1/2}, F_g = 3 \longrightarrow 6P_{3/2}, F_e = 4$ resonance to recycle the atoms. We will present additional details of these beams after introducing some of the other optical systems. A term diagram of the cesium D₂ line is shown in Fig. 2.2.

In the magneto-optic trap, it is the zero of the magnetic field that define the center of the trap. In the *absence* of the fields, the optical beams (cycling and repumping) that would otherwise form a MOT provide only a damping force without confinement. This damping force can effectively stop atomic motion for short times. This configuration is usually referred to as optical molasses. In a microscopic view, atoms in the optical molasses undergo a rapid series of spontaneous emission events, where each recoil is in an effectively random direction. This leads to three-dimensional diffusion in momentum. The competing process of damping from the optical molasses leads to a slow overall diffusion in atomic position. For the short times (~ 10 ms) that we image our atoms, it is a good approximation to say that the atoms are “frozen” in place, as we see in Fig. 2.3.

Elaborate methods have been developed for loading large numbers of atoms into a MOT. A notable example is the Zeeman slower for an atomic beam [Phillips82]. For our experiments, however, it has usually been sufficient to load the trap directly from background vapor in our vacuum chamber [Monroe90]. This is the case because most of our experiments have not required particularly large numbers of atoms. For some of the more recent experiments that involve state selection, (where the signal-to-noise ratio has been very low), we would actually benefit from much larger initial samples.

2.3 Optical Systems

A block diagram of the major optical systems is shown in Fig. 2.5. Two laser diode systems provide the light used in the MOT. The first of these is the molasses laser, which provides the light that is near the cycling transition. The other is the repump laser that returns atoms to the cycling transition. The lasers themselves will be described shortly (§2.3.1). A third laser, an argon-pumped Ti:sapphire laser, provides the light that forms the interaction potential. All of these optical systems (except the wavemeter) are installed on a Technical

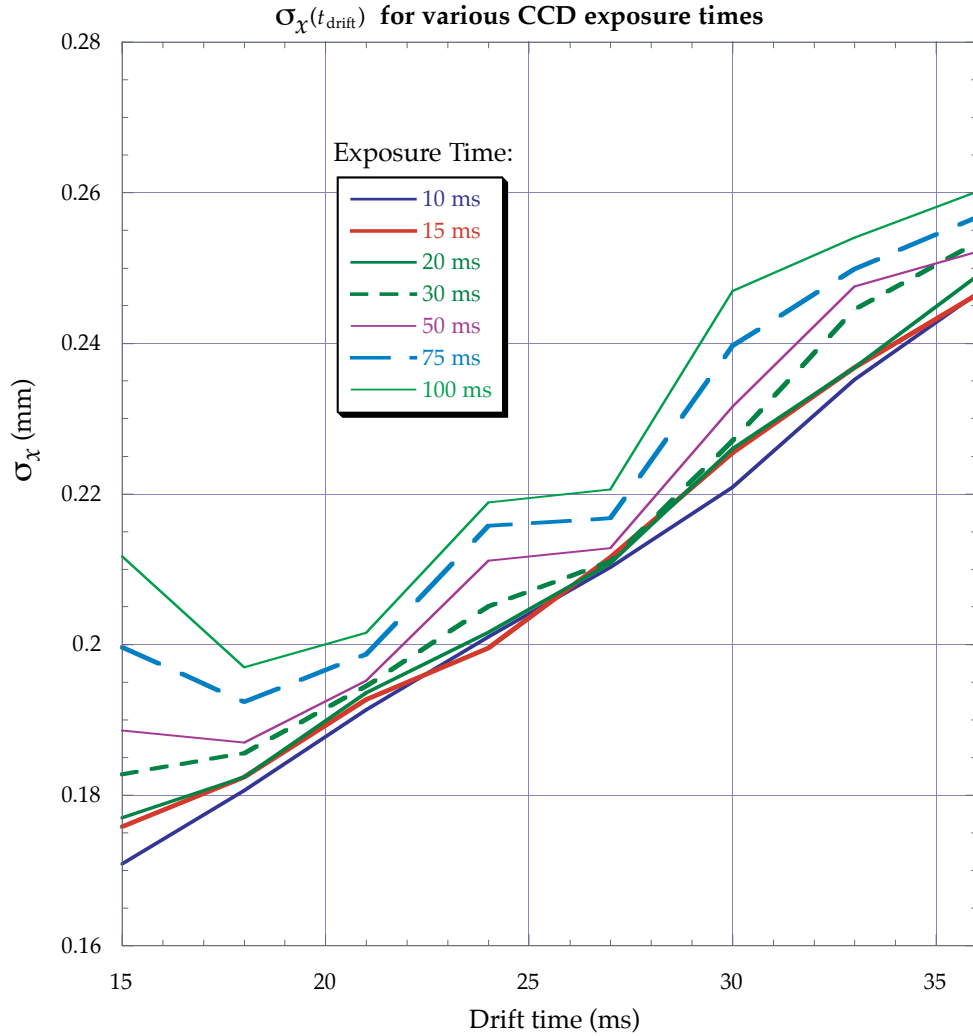


Figure 2.3: Characterization of freezing molasses. To collect this data, we loaded the MOT and released the atoms for a certain ballistic expansion (free drift) time t_{drift} . After the drift time, we turn on the freezing molasses and image the atomic distribution for a specified amount of time (10 – 100 ms). We then fit a Gaussian function to the one-dimensional position distribution to find the characteristic size σ_x . If the atoms are *truly* frozen in place by the optical molasses, the final size that we observe should not depend upon the exposure time. (If this is the case, then we can determine the momentum distribution from the curve.) As it turns out, significant diffusion in the molasses only occurs for drift times larger than about 20 ms. In most of our experiments, the total exposure time has been between 10 and 20 ms.

Manufacturing Corporation 4' × 12' damped optical table, which has been described by several visitors to our laboratory as one of the most complicated optical benches in the




Figure 2.4: Photograph of a region of the optical table. In this view we are looking directly at the vacuum chamber, which is obscured by hundreds of optical mounts. This picture illustrates how difficult it can be to visualize the optical systems by directly looking at them. A (more useful) block diagram of the optical systems is shown in Fig. 2.5.

world (see Fig. 2.4). The support electronics for the optics are kept below the optical table, on a platform suspended above the table, and in equipment racks separate from the table. This equipment includes the drivers for our many acousto-optic and electro-optic modulators, laser power supplies, servo-locking electronics, and assorted equipment involved in the computerized control of the experiment.

Let us first trace the beam path of the molasses laser. The beams that leave both diode lasers have an elliptical cross section. The aspect ratio is corrected to be nearly circular by Melles Griot anamorphic prism pairs (not shown) that are placed just after each laser. After this prism pair, the molasses laser beam passes through two Conoptics model 713 Faraday-effect optical isolators in series. The isolators have been measured to yield 40 dB and 37 dB of isolation, respectively. This double layer of isolation is necessary because the laser is extremely sensitive to optical feedback. Even with the two isolators, each of the downstream retroreflectors (at the MOT) must be slightly misaligned in order to prevent

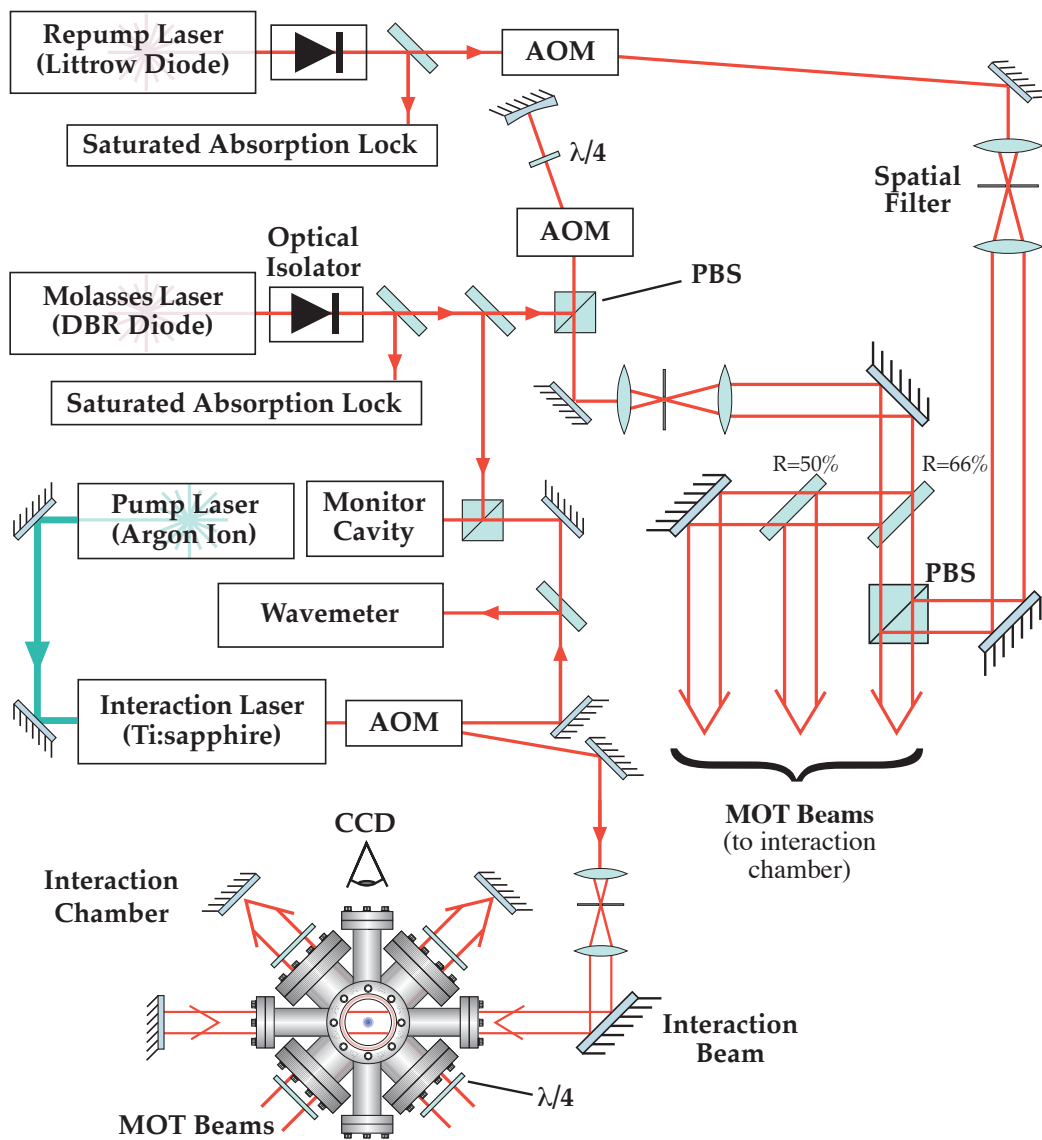


Figure 2.5: Block diagram of the essential optical systems. Two diode lasers generate the light that is used for trapping and cooling our cesium sample. An argon-pumped Ti:sapphire laser creates our one-dimensional interaction potential. Polarizing beam splitter (PBS) cubes are used to combine beams and double-pass the acousto-optic modulators (AOMs). Quarter-wave plates are indicated by the designation $\lambda/4$. The “MOT” beams indicated here are the combined repump and optical molasses beams. These three beams are retro-reflected through the interaction chamber along three orthogonal axes.

feedback from destabilizing the laser. About 75 mW of light is available after the isolators.

A weak cylindrical telescope at this point (not shown) is used to correct some residual

astigmatism in the beam. A small fraction (10%) of the light is then split off and directed to an FM saturated absorption locking setup (§2.3.2). The beam is then double-passed through a tunable acousto-optic modulator (AOM) centered at 80 MHz. This AOM is an IntraAction Corp. model ATD-801AL2, specified to operate between 60 and 100 MHz and is controlled via a driver from IntraAction. The laser itself is locked 195 MHz to the red of the $F_g = 4 \rightarrow F_e = 5$ cycling transition. Using the double-passed AOM, we can then vary the output detuning between 75 MHz red and 5 MHz blue of the cycling transition. For normal operation of the MOT, we keep the output detuning near -15 MHz (red). After shifting the frequency, we prepare the beam to enter the interaction chamber. The beam is spatially filtered by focusing it through a $25 \mu\text{m}$ pinhole and subsequently collimated with a beam waist of about $w_0 = 11$ mm. We typically have 27 mW of power immediately after the pinhole. The beam is then split into three parts of roughly equal intensity with two 2" diameter beamsplitters of reflectivity $R = 66\%$ and $R = 50\%$, respectively. One of the three components (the vertical beam) is combined with the repumping beam using a 2" polarizing beamsplitter (PBS) cube. The beams each pass through a 1.5" diameter quarter-wave plate to circularly polarize the light before it enters the interaction chamber. The three beams enter the chamber in three orthogonal directions. One beam travels vertically and the other two are at a right angle in a plane parallel to the optical table. A second set of quarter-wave plates follows the chamber, and the beams are retroreflected through these to realize the $\sigma^+ - \sigma^-$ optical molasses. The wave plates and retroreflectors are mounted on structures attached directly to the vacuum chamber flanges.

The repump beam path is slightly simpler. After the anamorphic prism pair, the beam passes through a single optical isolator like those used for the molasses beam. 10% of this beam is split off to an independent saturation absorption locking setup. The repump laser is locked 100.5 MHz to the red of the $F_g = 3 \rightarrow F_e = 4$ resonance. A fixed frequency 100.5 MHz AOM, also from IntraAction, shifts the output directly onto resonance. The light is spatially filtered and collimated to a size close to that of the molasses beams. The typical output power is 16 mW after spatial filtering. The beam is combined with the vertical MOT

beam and is retroreflected through the chamber.

The interaction beam is derived from the Ti:sapphire laser output beam. It is picked off with a fast, fixed-frequency 80 MHz AOM, IntraAction model ATM-801A2-2. This beam is spatially filtered with a 50 μm pinhole and collimated with a beam waist of about 1.5 mm. The beam is periscoped up to the level of the chamber and retroreflected through the chamber. In later experiments, a linear polarizer was placed after the periscope to ensure that the polarization of the beam was as linear as possible. Unlike the diode laser beams, the total output power available after the spatial filter varies with the power of the Ti:sapphire laser. This power tends to have significant ($\sim 50\%$) variation on a timescale of months to years, depending on many factors including the performance of its argon-ion pump laser. Notwithstanding this, a typical (acceptable) output power is 390 mW after the spatial filter. The interaction beam enters the chamber in the plane of the horizontal optical molasses beams, at a 45° angle to them (as viewed from above). The CCD camera is oriented perpendicularly to the standing wave so that it can best observe the atomic distribution along the axis of the standing wave.

The part of the Ti:sapphire beam that is not picked off by the interaction beam AOM (the 0th order component) proceeds to frequency-monitoring components (§2.3.2). Most of the power is directed through an optical fiber to the wavemeter, which is located on an adjacent optical table. The rest is directed into a Fabry-Perot monitor cavity, along with some of the light from the molasses laser. This setup is used to monitor frequency drifts of the Ti:sapphire laser relative to the molasses laser.

It is important to note that the AOMs on these three beams serve as optical *switches*, in addition to their frequency shifting roles. There is typically some “leakage” of radio-frequency (RF) power from the AOM drivers to the modulators, even when the control signal applied to the driver input is zero. This can, in some cases, lead to leakage of laser light into the interaction chamber when it is not intended. For the molasses beam, which is double-passed, this is not an issue, because the small amount of leakage is strongly attenuated by the response function of the AOM. On the repump beam, an RF switch was

placed in the AOM driver, between the frequency-synthesizing elements and the RF power amplifier. This switch is used to provide an additional layer of attenuation when the repump laser beam is intended to be off. It is particularly important to avoid any leakage of the diode lasers because they are so close to resonance and consequently, very low levels of leakage can potentially alter the internal atomic state. There is a very low level of leakage (up to 0.1%) through the interaction beam AOM, although the exact cause of this leakage has not been determined. (An accidental reflection is suspected.) Because the interaction beam is much further detuned from resonance, we do not expect this to have a significant effect.

2.3.1 Optical sources

An advantage of the particular cesium transition that we are using is the availability of diode lasers at this wavelength. Diode lasers have the usual array of advantages and disadvantages as compared to other laser sources. Most of the higher power devices that are available are in the form of “broad stripe” diodes that have rather poor spatial mode quality, and may be difficult to lock to an atomic transition. In the lower-power devices that we employ (up to 150 mW), the overall cost of operation is expected to be low compared to that of other types of laser systems.

The molasses laser is a distributed Bragg reflector (DBR) diode laser, SDL, Inc. model SDL-5712-H1. A DBR laser is produced with a frequency-selective structure and typically produces single frequency output. This laser is specified to produce 100 mW of output power at its maximum operating current of 180 mA. The laser is mounted in a TO-3 package with integral thermoelectric cooler (TEC) and thermistor. The laser is mounted on a passive (air-cooled) heatsink manufactured by SDL (model SDL-800). The collimating lens is an antireflection (AR) coated 5 mm lens manufactured by Rodenstock. It was originally mounted on a Line Tool XYZ translation stage, however this proved to exhibit gradual creep and needed weekly adjustment. We later glued the collimating lens directly to a structure attached to the front plate of the heatsink [Steck01b]. After gradually creeping for a few




Figure 2.6: Photograph of the DBR diode laser. This laser provides the light that forms the optical molasses in our experiment. The diode itself is housed in the gold-tone TO-3 package (A). The package is mounted through an adaptor plate to an air-cooled heatsink (B). A special structure (C) is mounted to the front of the adaptor plate. This structure holds the collimating lens tube (D) on the top of two parallel glass rods. The entire structure is clamped to the optical table (directly below the heatsink). The housing for the repump laser is visible in the background, as are several racks of control electronics. (Photograph: D. Steck)

weeks, the lens position settled in, and it has not required adjustment for at least several months. A particular disadvantage of the DBR laser that we are using is that it is no longer manufactured by SDL.

The repump laser is a grating-stabilized diode laser in the Littrow configuration [MacAdam92]. In this configuration, the collimated output of the laser is incident upon a diffraction grating that reflects some of the light (the first order) back into the diode itself, effectively forming a standing-wave laser cavity bounded by the grating and the rear facet of the diode chip. When properly aligned, the grating is frequency-selective and encourages the cavity to operate with a single frequency.

The laser diode is SDL model 5421-G1, which is housed in a standard 9 mm G1




Figure 2.7: Photograph of the repump laser, a Littrow configuration grating-stabilized laser. The laser diode itself is housed in a collimation tube that is mounted in the aluminum bronze structure. The wires that lead to the diode are visible near the center of the photograph (A). The angle of the diffraction grating (B) is adjusted coarsely with a 1/4-80 set screw (C). Fine control over the angle is made possible by a stack of piezoelectric discs, which the ball-end of the set screw contacts. The fine pink and blue wires provide the high voltage to drive the piezo stack. The lucite cover of the structure serves to keep out dust, improve thermal stability, and reduce acoustic coupling. The beam exits the structure through a Brewster window (not visible) on the right side of the structure (D). A solid red arrow is used to illustrate the beam path, which is otherwise not visible. (Photograph: D. Steck)

package. This diode produces 150 mW of output power at its maximum operating current, 200 mA. The diode package is mounted in a Thorlabs LT230P-B collimation tube that includes a 4.5 mm focal length collimating lens. This tube and the diffraction grating are mounted in a flexure structure made of 954 aluminum bronze, an elastic material with high thermal conductivity. A photograph of the external-cavity structure is shown in Fig. 2.7. The flexibility of the metal itself is used to allow positioning of the grating, rather than moving parts or common hinges. The angle that the grating makes with the diode laser output is determined coarsely by the position of an 80-thread per inch (TPI) screw. The fine angle adjustment is executed with a stack of three 8 mm diameter by 2.54 mm piezoelectric

ceramic discs from American Piezo Ceramics. The stack is driven at up to 1 kV by one channel of our Trek model 601B-4 high voltage amplifier. The grating was purchased from Edmund Scientific and coated with a thin (500 Å) layer of gold to improve its reflectivity at our wavelength. The Littrow condition is met with an angle of 31° , where 67% of the incident light goes into the 0th order output beam. The laser can be continuously tuned (that is, without hopping in frequency) for approximately 1 GHz.

We are not actually able to operate the laser near its maximum operating current, however, because we have had difficulty obtaining stable operation at the atomic resonance with higher currents. This has not been a major concern because we do not need much repumping light. (Our 16 mW at the chamber are more than sufficient.) Typically, the laser is operated with an injection current near 100 mA, and the output of the external cavity laser is near 40 mW. The laser structure sits on a Melcor TEC that is connected to an aluminum plate that serves as a heat sink. A thermistor in the bronze structure provides feedback to the temperature controller. A lucite cover helps to reduce thermal variation and acoustic sensitivity. The general design of the external cavity laser is modeled after a design that made its way to our lab via Konstanz, Germany. Many more details about this laser system are available in the earlier dissertations [Klappauf98c; Steck01b].

Both diode lasers are controlled by NIST laser diode current controllers. Matching NIST temperature controllers stabilize the temperature of the thermistors in the two diode housings by controlling the current through the TECs. These controllers were designed and constructed at NIST in Boulder by the group of Leo Hollberg. The units are in the nuclear instrument module (NIM) form factor, and we operate them in a Canberra 1400 NIM crate. We have, generally speaking, had excellent luck with these controllers. We have not (over the course of the experiment thus far) blown out a single laser diode, which is an excellent credit to them. We have, on the other hand, had a somewhat mysterious interruption during longer data runs. The current controller has a tendency to suddenly switch the diode off. During our last data run on the chaos-assisted tunneling experiment, the two diodes each tended to shut off (on average) once per 24-hour period. Our best guess is that there is a

trigger-happy protection circuit in the current controllers, but it is better to err on the side of safety.

For the most part, the diode lasers are turnkey systems. Their operation usually involves an hour or so of warmup time. After that, they require a moment's work with their locking electronics to lock them to the proper frequencies. The optical setup for each laser is generally speaking quite stable, and we do not adjust any part of their beam paths on a regular basis. During long periods of operation, a common interruption is the unlocking of one or more of the frequency-lock systems. This can be caused by acoustic noise, human-induced vibration of the table, and/or drifts in the temperature of the laboratory over the course of several hours.

Besides the diode lasers, we also employ an argon-pumped, home built Titanium-sapphire (Ti:sapphire) ring laser. This laser was originally built by Patrick Morrow as a dye laser for earlier experiments in our laboratory [Fischer93]. The laser was converted to a Ti:sapphire laser by Bruce Klappauf and Daniel Steck, the year before I joined the experiment. Details of the conversion, the resulting laser design, and performance are described in the dissertations of Klappauf and Steck [Klappauf98c; Steck01b]. The laser typically produces about 480 mW of output power in a clean gaussian beam, although the power does again depend upon the pump laser performance. The Ti:sapphire laser is pumped by a Coherent Innova 90 argon-ion laser. This laser is specified to produce 5 W of multi-line continuous wave (cw) power. Depending upon the age of the tube, the laser actually outputs between 7 and 8 W of blue-green light, which is the only visible light on our optical table. The Ti:sapphire laser also is sensitive to the spatial mode quality of the pump beam, which tends to degrade as the tube ages.

The Ti:sapphire is much more difficult to operate on a daily basis than the diode lasers. We usually wait until the argon laser has been on for at least an hour before directing the beam into the Ti:sapphire laser. This also allows time for our chiller to cool the water that cools the Ti:sapphire crystal itself. This cooling water is recirculated through a Neslab chiller and is kept near 11°C, just *above* the point where condensation on the water tubing

(hoses) becomes a problem. This cooling water also cools the anti-Helmholtz coils and the CCD camera. After the warmup/cooldown period, we let the pump beam into the Ti:sapphire laser cavity by removing a beam block at the pump laser output. There are (quite literally) many knobs to turn when adjusting the laser. A substantial degree of both skill and patience is necessary to get the laser to the proper frequency at an acceptably high output power, while also ensuring that the system is *stable* in frequency and in power.

2.3.2 Laser stabilization and monitoring

For those laser beams that are near to an atomic resonance, the resonances themselves are used as frequency references. In the optical paths that we described earlier, some of the light from both the molasses and repump lasers is split off and directed to independent saturated absorption locking setups. Saturated-absorption spectroscopy is a means of performing precision spectroscopy on a sample of atoms in a room-temperature vapor cell [Pappas80; MacAdam92]. Normally, the spectrum of a room temperature gas cannot be resolved because the thermal motion of atoms leads to Doppler broadening on the order of 0.5 GHz [Demtröder96]. This value is comparable to the entire width of the cesium excited-state manifold ($6P_{3/2}, F = 2 - 5$), and makes it impossible to resolve the individual spectral lines. Saturated absorption spectroscopy uses counterpropagating beams through a vapor cell to resolve only those atoms that are at (for example) zero-velocity, and thus simultaneously resonant with the beams from both sides of the cell. This type of spectroscopy is a standard technique and yields an atomic spectrum of sufficient strength and resolution to serve as a reference for our lasers. Peaks in the saturated absorption spectrum occur at resonances between hyperfine states, but also at *crossover* resonances, located at each midpoint between two hyperfine resonances. Each of these signals appears as a perturbation (a “Lamb dip”) in the Doppler-broadened spectrum. In our locking schemes, we lock each laser to the midpoint of a crossover resonance. One reason for this is that these particular resonances are the strongest features in their respective spectra.

The repump laser’s locking setup is fairly straightforward, so we will review it first.

The locking optics were originally located right next to the laser itself, adjacent to the interaction chamber. In order to free up some table space for the 3-D far-detuned lattice that we installed later, the setup was eventually moved to a secondary optical breadboard (“the second floor”) located directly *above* the argon-ion laser. The heart of the locking setup is a small (75 mm long) glass vapor cell from Environmental Optical Sensors, Inc. (EOSI) that is maintained at room temperature. Using a 3/8” thick uncoated optical flat, we pick off two weak, parallel beams that are transmitted through the vapor cell. One of these serves as the “probe” beam for the saturated absorption spectrum, and one serves as a reference beam. The majority of the beam is not reflected from the optical flat, and is directed as a “pump” beam to the other side of the vapor cell. This beam is directed to overlap the counterpropagating probe beam in the vapor cell. The probe and reference beams are focused onto a pair of photodiodes that operate in differential mode. The intensity of the reference beam gives the Doppler-broadened lineshape of the vapor cell, and the probe beam gives the raw saturated-absorption spectrum. Subtracting away the reference signal yields an output spectrum that shows only the features of the resonances themselves.

We would like to lock the laser to the center of a resonance, and it is helpful to have a dispersive error signal that we can use to lock the laser. We modulate the laser frequency directly by feeding a weak modulation to the high-voltage amplifier that controls the grating position. The modulation frequency is typically 12 kHz. We then process the subtracted photodiode signal with an EG&G 5204 lock-in amplifier to create a dispersive error signal. This error signal (shown in Fig. 2.8) is in return fed back to a home-built lock box that is mounted in our NIM crate. The lock box uses proportional and integral feedback to steer the grating position as necessary to keep the laser on resonance. The direct modulation of the output laser frequency is generally undesirable. However, the total deviation that the modulation induces is of order several MHz (by inspection of the modulated spectrum). This type of dither should not cause any problems because atoms in the MOT spend nearly all of their time in the cycling transition where they do not “see” the repump light. The specific transition that we lock the laser to is the $F_g = 3 \rightarrow F_e = 3, 4$ crossover resonance.




Figure 2.8: Saturated absorption spectrum of the repump laser. Rather than the direct absorption signal, we display the output of the lock-in amplifier that produces the dispersive lineshapes that we use for locking the laser. We recorded this curve with our digitizing oscilloscope as the laser was swept across resonance many times. This curve was averaged over many sweeps to improve the signal to noise ratio. The range of the sweep is limited to the $F_g = 3 \rightarrow F_e$ transitions, and each spectral line is labeled by the corresponding value of F_e , or the two values of the crossover lines. In our experiment, we lock the laser to the $F_g = 3 \rightarrow F_e = 3, 4$ crossover resonance. Not coincidentally, this is the strongest feature in the spectrum.

From Fig. 2.2, we see that the splitting between the two lines is about 201 MHz, so the crossover resonance is close to 100.5 MHz to the red of the repumping ($F_g = 3 \rightarrow F_e = 4$) transition.

The molasses laser is locked by an FM saturated absorption setup that is considerably more complicated than that used for the repump laser [Bjorklund83]. The light for the locking setup is again picked off from the main beam immediately after the optical isolation stage. The light, about 8 mW total, is first spatially filtered with a 75 μm pinhole. The spatial mode is important because the modulators that we use in this setup are sensitive to the mode quality. The light is then split into pump and probe components. The pump portion is double-passed through a chopped 70 MHz AOM to pulse the beam and shift the

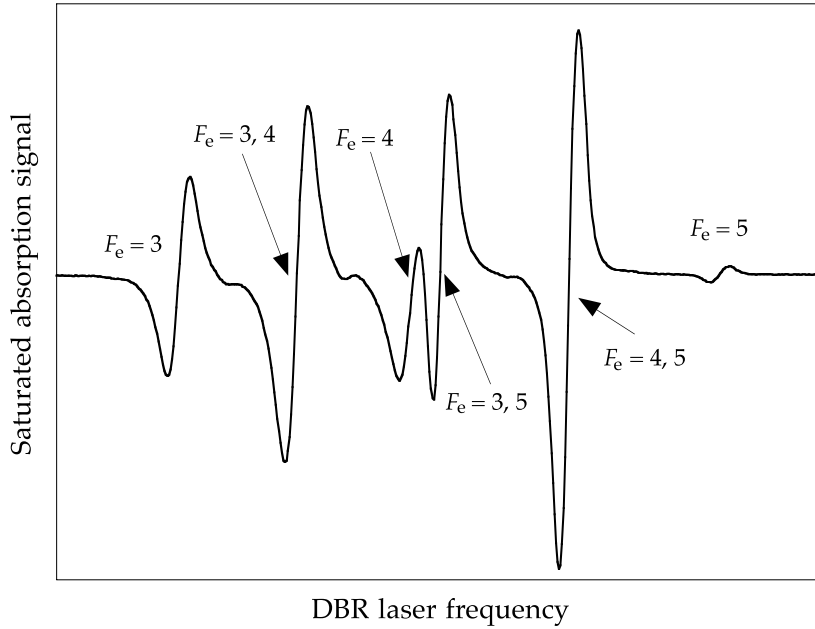


Figure 2.9: Saturated absorption spectrum of the molasses laser. This error signal is derived from the FM locking setup and was averaged over many sweeps to improve the signal to noise ratio. The range of the sweep is limited to the $F_g = 4 \rightarrow F_e$ transitions, and each spectral line is labeled by the corresponding value of F_e , or the two values of the crossover lines. In our experiment, we lock the laser to the $F_g = 4 \rightarrow F_e = 4, 5$ crossover resonance.

frequency up by 140 MHz. The power to the AOM is chopped with a 50 kHz square wave. This pump beam then enters an EOSI vapor cell like that used in the repump frequency lock. The probe beam enters from the opposite end of the vapor cell and overlaps the pump beam. Before it enters the cell, however, it passes through a ConOptics model 350-40 electro-optic phase modulator. This modulator is driven at about 11 MHz to realize the necessary dither in the laser frequency. The probe beam intensity is monitored after the vapor cell with a low noise photodiode from New Focus, model 1801. The signal from the photodiode is mixed with the 11 MHz reference signal, and this signal is in turn fed into a Stanford Research Systems SR510 lock-in amplifier, referenced to the 50 kHz chopping signal. Our phase-sensitive detection of the phase-modulated signal yields a dispersive line-shape, and locking in to the chopped pump signal makes the output signal insensitive to the Doppler-broadened background. An error signal derived from the lock-in amplifier is used

to steer the frequency of the DBR laser, which is easily manipulated by directly changing the injection current. A second “lock box” like that used on the repump laser contains proportional-integral feedback loops to keep the laser on resonance.

In the locking setup, we have shifted the frequency of the pump beam by 140 MHz. The pump and probe beam then interact with a velocity class in the vapor cell that is Doppler-shifted by 70 MHz relative to the lab frame. (A 70 MHz shift is well within the Doppler line width of the spectrum.) A second effect of the frequency shift is that when the output signal is centered on a specific spectral feature, the laser itself is at a frequency 70 MHz to the red of the spectral feature. The specific feature that we lock to is the $F_g = 4 \rightarrow F_e = 4, 5$ crossover resonance which is roughly 125 MHz to the red of the cycling ($F_g = 4 \rightarrow F_e = 5$) transition. The molasses laser output is thus locked to a frequency of 195 MHz below the cycling transition, as we had claimed earlier. Despite its complexity, this setup does have advantages over the lock used for the repump laser. First, note that the laser output itself is *not* modulated. This is important because the performance of optical molasses and the MOT can be fairly sensitive to the detuning of the laser. Secondly, this initial detuning allows us to have a wide frequency-tuning capability. As we have mentioned, the molasses light is double-passed through a tunable 60 – 100 MHz AOM that lets us vary the final detuning from slightly blue of resonance to 75 MHz red. Again, since the optical molasses laser is sensitive to the detuning, this detuning capability gives us some degree of freedom in the behavior of the MOT.

Additional details of both diode laser locking schemes, including diagrams of the optical and electronic systems that we have described here are presented in Klappauf’s dissertation [Klappauf98c]. These systems have provided reliable operation for several years, and the only major change since this documentation has been the relocation of the repumping lock optics.

Since the beams from the Ti:sapphire laser create the interaction potentials that we use to do the “physics” part of our experiments, the performance of these beams is critical. We are concerned with the frequency, intensity and phase stability of our optical lattice.

Typically we operate in a regime where we are very far detuned from resonance, as compared to the diode lasers. The typical detuning of the Ti:sapphire laser for our experiments has been in the range of -6 to -40 GHz, where this detuning is always specified with respect to the $F_g = 4 \rightarrow F_e = 5$ cycling transition. Because of the magnitude of this detuning, it is not presently feasible for us to directly lock the laser to one of the atomic transitions. Originally, this laser had the capability to lock to a dispersive error signal from an Invar reference cavity with a 1.5 GHz free spectral range (FSR) [Hänsch80]. As a practical matter, we have found that the frequency stability was not substantially improved by actively locking. The stability of the free-running laser is typically adequate for our purposes except for occasional “hops” that occur simultaneously with changes in intensity. Mode hopping may be due to gradual changes in (for example) the ambient temperature in the laboratory and tends to occur over a period of hours or days. The laser is most stable, generally speaking, late at night and on weekends, particularly if it has been running continuously for more than a day. These failures tend to be one of the rate-determining factors in data collection, and cannot be easily prevented.

The absolute frequency of the Ti:sapphire laser is coarsely determined by the wavemeter shown in the block diagram (Fig. 2.5), which is the only optical system that is not on the main table. The wavemeter is a NIST LM-10 Lambda Meter, which is a traveling-arm scanning Michelson interferometer [Hall76]. It works by comparing the wavelength of our “unknown” beam to the known wavelength of a temperature-stabilized helium-neon (HeNe) laser. We have had to replace the tube in the reference laser at least twice, because it is the nature of gas lasers tubes to eventually burn out. When this has happened, it has resulted in downtime for the experiment as a whole, because we use the wavemeter on a near-hourly basis to monitor the laser frequency. The wavemeter’s traveling arm consists of an “air car,” a one-dimensional air-bearing translation stage, that carries a pair of cornercube reflectors. For most of the time that we ran this experiment, we floated the air car with dry nitrogen from a compressed gas cylinder. In the past year, however, it has run off of a compressed air system that was installed in our laboratory by Patrick Bloom, an undergraduate research

assistant. Each reading is made by counting the ratio of Ti:sapphire laser fringes to HeNe laser fringes from the interferometer while the air car moves along its full range of travel. We use a Hewlett-Packard 5334A counter to count both sets of fringes and take their ratio. The fringe counting is gated with an averaging time of 7 s, which corresponds to the one-way travel time of the air car. Typically, we can use this instrument to determine the output frequency of our laser to within about 500 MHz. It is worth noting that the Ti:sapphire beam that heads to the wavemeter is only intense enough to make a wavemeter reading when it is not picked off by the AOM to become the interaction beam. Because of the regular interruptions to this light during experiments and the need for averaging, we are only able to use the wavemeter *between* individual data runs.

In order to determine the wavelength of the laser on a more precise scale, we employ a scanning confocal Fabry-Perot interferometer monitor cavity [Demtröder96]. This cavity is one of a family that I built as my first project on the experiment, back in 1997. It has two spherical mirrors of 99% reflectivity at 852 nm, separated by their radius of curvature, 5 cm. The front (input) mirror is mounted in a kinematic stage adjustable with three New Focus 6-80 miniature fine adjustment screws. The rear mirror is not adjustable in its angle, but sits on a short stack of piezoelectric ceramic rings. The stack gives the rear mirror a translational degree of freedom along its axis of symmetry that allows us to sweep the total spacing between the mirrors. The cavity has a free spectral range of 1.5 GHz and a finesse of about 300. Since we always operate this cavity in a scanning mode of greater than one FSR, absolute thermal drifts are not important. However, it is best that they be minimized over a short time scale. The cavity structure consists of a heavy brass spacer between the mirror mounts (for thermal mass) and an anodized aluminum housing that minimizes air currents through the structure. The structure also houses a photodiode that monitors the intensity of light leaving the cavity. Typically, we scan this laser over roughly two free spectral ranges using a triangle-wave with frequency on the order of 100 Hz that is amplified to high voltage by an independent channel of the Trek amplifier.

The optical input to the monitor cavity consists of combined beams from the

Ti:sapphire laser and the molasses laser. The molasses laser is always locked to the same frequency, 195 MHz to the red of the cycling transition. If we scan the frequency of the Ti:sapphire laser until the two peaks on the output of the monitor cavity coincide, then we know that the frequency difference between the two beams is an integer number of free spectral ranges, $\Delta\nu = n \cdot 1.5$ GHz, for some integer n . The Ti:sapphire detuning from the cycling transition is then -195 MHz + $n \cdot 1.5$ GHz. Noting that the interaction beam is upshifted by 80 MHz by the AOM that acts as its shutter, the detuning of the interaction beam from is then $\Delta_L = -115$ MHz + $n \cdot 1.5$ GHz.

Operationally, we use the wavemeter to get the Ti:Sapphire laser to within about 1 GHz of the desired detuning. We then adjust the detuning to match up the peaks in the monitor cavity. Once these signals have been aligned, their gradual separation is visually monitored on an oscilloscope screen. This method has proven to be quite reliable. The typical drifts in frequency are on the order of 100 MHz per hour, although the laser can be more stable during long runs, as we have described.

Besides frequency stability, the intensity stability of the optical lattice is also important. When operating properly, the rms amplitude noise on the interaction beam has been measured to be on the order of 0.2%, with both long-time and short-time integration. In this case, short term means over the 2 ms timescale of individual experiments, and long term means over the course of seconds between subsequent experiments. Over the short term, the noise turns out to be dominated by fluctuations at 2.4 kHz which are introduced by active components in the Ti:sapphire laser itself. Over longer time scales, the total output power tends to vary substantially more, primarily in slow drifts that we attribute to thermal variation in the laboratory. In general, the slow drifts in power exhibit similar behavior to the slow drifts in frequency that we have discussed earlier. Similarly, rapid hops in the laser intensity coincide with jumps in the frequency.

We commonly monitor the power level of the Ti:sapphire laser during the course of experiments to ensure that no substantial drifts or hops have taken place. Again, part of the reason for this is that the frequency and intensity changes tend to be correlated.

The computer that controls the experiment routinely monitors a photodiode that is built into the Ti:sapphire laser. This is used to provide a simple assurance that the laser system is functioning properly. We are, of course, concerned with the intensity variations in the lattice when it gets to the interaction chamber. Monitoring the laser itself does not account for possible variation in the diffraction efficiency of the AOM or minor optical misalignment that may result in reduced transmission through the spatial filter. One of the “mirrors” in the interaction beam path (after the spatial filter) is actually a beamsplitter with 99% reflectivity. The light transmitted through this beamsplitter falls on a fast (50 MHz) photodiode from Thorlabs. We use this detector, read out by a Tektronix TDS 524A 500 MHz digitizing oscilloscope, to monitor the intensity of the interaction beam during the individual experiments. For example, we have used it to digitize the pulse profile of individual “kicks” in our kicked-rotor type experiments.

A third type of stability that we require is in the phase of the optical lattice. The stability of the standing wave is determined by the position of the retroreflecting mirror. This stability is critically important as we expect the quantum effects that we are studying to be very sensitive to phase noise. The retroreflecting mirror is held in a mount that is rigidly clamped to the vacuum chamber flange. Because of the design of the chamber base, the mirror position is quite stable against motion along the axis of the standing wave. The phase stability was measured using a Michelson interferometer to have typical phase fluctuations of about 8% of the lattice period over the 2 ms time scale of these experiments. Since the reference components of this interferometer setup are presumed to be less stable than this mirror, this phase noise measurement is likely to be a significant overestimate. We have worked hard to remove any potential sources of noise on the optical table. The only moving parts on the optical table during the experimental sequence are those in the CCD camera shutter assembly, and the shutter only begins to move after the end of the interaction. The remaining dominant source of vibration is that of the cooling water flowing through the argon-ion laser. The cooling water that runs through the anti-Helmholtz coils is in mechanical contact with the chamber, however, this flow is much slower and does not

appear to introduce appreciable vibration to the optical table.

2.3.3 Optical lattice calibration

Beyond simply ensuring that the far-detuned optical lattice properties are consistent, it is also important to be able to determine its properties absolutely. We have already discussed the procedures by which we determine the absolute frequency of the interaction beam. Another critical quantity that we must determine is the *intensity* of the beam, which is necessary to find the total strength of the interaction with the atoms. Generally speaking, intensity calibration consists of knowing two things: the total power of the beam and the Gaussian spot size.

The total power is measured with an optical power meter with a detector placed after the spatial filter in the interaction beam. Our primary power meter is a Newport model 1825C power/energy meter. This meter reads out the signal from a Newport model 818-SL semiconductor detector head that incorporates a 1000:1 attenuator, model 883-SL. Both the meter and the detector head are independently calibrated by Newport. The calibration certificates for the power meter and detector each specify an absolute uncertainty of $\pm 1\%$, although we do not actually trust them to this level of accuracy. We also have a calibrated thermopile detector from Newport and a Coherent Fieldmaster power meter that is also based upon a semiconductor detector. While each of these meters claims to be accurate to within a similar range of uncertainty, they typically disagree with each other by slightly more than this amount. We have, for the sake of consistency, generally relied upon the measurements of the 1825C/818-SL system. Typically, this sort of absolute power measurement overestimates the power in the interaction region by about 8%, because it does not account for losses in the beam path after the point where the measurement takes place. The losses include a tiny fraction at each optical surface, including residual reflection from the interaction chamber windows.

The intensity in the interaction region is determined not only by the total power, but also by the cross-sectional area of the beam. We have several procedures for measuring

this area, which is quantified by a beam-waist parameter since the beam is Gaussian and well collimated. Our most basic method of measuring the beam size consists of scanning a knife edge transversely across the beam profile. The knife itself is a simple single-edge razor blade that we attach to a one-dimensional translation stage with micrometer readout. The translation stage is located just after the spatial filter, at the same location that we measure the total power. The power that is not blocked by the knife edge is essentially a one-dimensional integral over that part of the Gaussian profile, yielding an error function as a function of position. Typically we step the blade across the profile in $50 \mu\text{m}$ intervals while monitoring the “transmitted” power with the Newport power meter. A rear-panel analog output from the power meter is in turn read out by one of our data taking computers. This semi-automated procedure (the translation stage is advanced by hand) collects the array of intensity values and the computer fits this array to an error function to output the beam waist. With an independent knife-edge scanning in the other (transverse) direction, we find the beam waist along both the horizontal and vertical axes.

A secondary method that we use to measure the spot size involves the use of a charge-coupled device (CCD) camera. We use a kinematic stage inserted near the knife edge setup to redirect the beam at this point. The redirected beam is heavily attenuated and directed onto the CCD chip of a Sony XC-77 camera. Strong attenuation is necessary because the CCD is easily saturated. The camera has a $2/3$ ” (diagonal) CCD chip, and was chosen primarily for its large total area, which allows most of the interaction beam to fit on the chip. The output is a standard grayscale NTSC signal that we digitize on the computer using a Digital Vision, Inc. ComputerEyes/RT SCSI frame grabber. The images are imported to an analysis program that integrates the pictures along the horizontal and vertical axes. These integrated profiles are fit directly to Gaussian functions to reveal the spot size. This method is not without flaws, however. The CCD images always display interference fringes that we believe arise from reflections off of the CCD chip itself and the window that covers the chip. To some degree these fringes can be integrated out, and the curve fits seem to be fairly robust against the fringes. A secondary problem is that

the absolute value of the beam waist (as measured this way) depends slightly upon the total intensity, presumably as a result of camera saturation effects. Nonetheless, the images from the CCD camera are very useful for making sure that the beam is close to circular in cross-section.

The knife-edge method is precise, however it is time consuming and we can only measure the beam on the way *to* the interaction chamber. The beam is retroreflected through the chamber, so a good way to verify that the beam is properly collimated is to verify that the beam size is the same on its way both to and from the interaction chamber. The location where we can measure the spot size is about 2 m “upstream” in the beam path from the interaction chamber, and the round trip optical path length from that point through the chamber and back is about 5 m. We can insert at this point a secondary kinematic stage that incorporates a pellicle beamsplitter to direct the beam on the way *from* the interaction chamber to the CCD camera setup. The pellicle is thin enough that it does not significantly deflect the beam on the way to the chamber, and on the way back, it reflects some light into the CCD. Using the camera, we can then measure the ellipticity and spot size of the beam at point effectively separated by 5 m in order to ensure that the beam is properly collimated when it is in the interaction chamber.

It is important to note that these measurements of the total power, beam profile, and collimation are all invasive calibration procedures— we cannot perform them during data collection. With the total power measurements and the spot size measurements, we are usually able to determine the total intensity in the interaction region, during an experiment only to within about $\pm 10\%$. However, in experiments where it is critical to know the absolute intensity to higher precision, we have relied instead upon *physical* calibration methods where atomic motion is sensitive to the intensity. We will come back to these methods later, but an example calibration is to experimentally trace the $D(K)$ curve in the kicked rotor that we described in the introduction.

A secondary issue to determining the absolute intensity is that of keeping it constant. We have discussed some of the issues in long and short term intensity fluctuations in the

interaction beam. We have developed several methods of taking out the effects of both long and short-term drifts in the Ti:Sapphire output intensity. We can, again, directly measure the intensity of the interaction beam by using the Thorlabs photodiode after the spatial filter. In the kicked rotor experiments we typically digitize the entire pulse train and average them together to find the average intensity of each pulse. In cases where we wish to measure some parameter as a function of the kicking strength K , this gives us a number proportional to the instantaneous value of K that we can use to remove short-term drifts from our data. In some of the most recent sets of experiments, we have implemented a similar system to remove the effects of long term drifts (over a period of weeks) in the laser system. In a procedure that we call the “virtual power lock” (VPL) we measure the total available interaction beam power prior to each data run, typically several times per hour. We then *scale* the control signals that control the AOM(s) that pick light off from the Ti:Sapphire beam. Additionally, during these runs, we discarded individual data sets in which the total power drifted by more than $\pm 1\%$. Using this method, we can keep the total interaction strength very constant despite large (up to 20%) total variations in the Ti:Sapphire laser output power.

2.4 Interaction Chamber

All of our experiments take place in an ultra-high vacuum (UHV) system. The system can essentially be divided into two parts: the main interaction chamber, and a support system that is attached to it. Optical access to the chamber is critically important because of the many beams that end up in the chamber. Magnetic field control is also very important because a strong field is needed for the MOT and because stray fields can disrupt our experiments in other ways. A photograph of the chamber, the magnetic field coils, and some of the optical systems is shown in Fig. 2.10.




Figure 2.10: Photograph of the interaction chamber. The chamber itself (A) is circumscribed by anti-Helmholtz coils above (B) and below the chamber. The loops of magnet wire are visible as reddish-brown stripes, and it is not easy to see the lower coil. The chamber is further surrounded by three orthogonal pairs of nearly-square Helmholtz coils (C). The large windows of the chamber, such as the one visible between the circles (A) and (B), provide optical access for the many laser beams that enter the chamber. One of the optical molasses beams enters the chamber through this port after reflecting from a mirror in the mount in the foreground (D). The barrel of the CCD camera lens (E) is barely visible in this photograph. The camera points through the chamber at the flange covered by the red cap, to the right of circle (F). Two valves at that end of the chamber (F) are used to seal off the main chamber from the cesium source (not visible). The optics on the platform on the lower right hand side of this picture are those illustrated in Fig. 4.10.

2.4.1 Vacuum system

Let us first describe the vacuum system itself. The main interaction chamber is a custom, stainless steel 10-way cross made by HPS (a division of MKS Instruments). Eight of the “arms” are in the horizontal plane (parallel to the plane of the optical table). Four of these equally spaced radial arms end at 4.25” (diameter) ConFlat (CF) flanges, and alternate with four that end at 2.75” CF flanges. The vertical arms each have a 4.25” CF flange. All but one of these ports is sealed with a Kovar-sealed viewport. The viewports are AR coated on both sides for less than 0.25% reflection per surface, nominally, at our wavelength. The six




Figure 2.11: Illustration of optical access to the interaction chamber. In this cutaway view, the optical molasses beams (red) intersect at the center of the interaction chamber, where the atoms (yellow sphere) become trapped. These six beams enter through the six 4.25" flange viewports. The beams are steered by mirror mounts attached directly to four rods that are screwed into the viewports. Similar mounts (not shown) are used to hold quarter wave plates between the mirrors and the chamber windows. The two flanges on the left and right sides of the chamber (in this view) are those used for the interaction beam standing wave. The port in the foreground is that which the CCD camera looks through.

larger viewports provide optical access for the MOT beams, as illustrated in Fig. 2.11. Two of the smaller viewports (in opposition) are used for the interaction beam (as in Fig. 2.5). The CCD camera that we use for imaging looks through the remaining window, and the remaining flange (opposite the camera) is the interface to the rest of the vacuum system. The interaction chamber sits on a 5/8" thick base plate that supports its weight and holds it high enough above the table that we have optical access to the bottom.

The rest of the vacuum system is fairly simple and supports the essential goals of keeping the residual pressure low and allowing the introduction of cesium atoms. Initially, the vacuum system was pumped down from atmospheric pressure by a turbomolecular vacuum pump, backed by a mechanical roughing pump. The vacuum chamber was baked out at up to 220°C over the course of several days to improve the ultimate pressure. After the bakeout, our vacuum chamber was pinched off from the mechanical pumps, and we

have not broken vacuum since that time, which was in the spring of 1997. The vacuum system incorporates a Varian 20 l/s ion pump, model 919-0236, to maintain the vacuum. The ion current provides one measure of the vacuum pressure, and the system also includes an HPS model 10000 5836 Bayard-Alpert type ionization vacuum gauge. We have had some difficulty in obtaining reliable readings from the vacuum gauge, which we attribute to cesium contamination. The ion current from the pump corresponds to a pressure of about $8 \cdot 10^{-8}$ torr, but it appears that this may also be inaccurate for the same reason. The decay time (lifetime) of cesium atoms in the optical molasses is of order 4 s, which suggests that the background pressure is closer to $2 \cdot 10^{-9}$ torr [Monroe92; Steck01b].

Cesium vapor is introduced to our chamber from a 1 g metallic sample in a glass break-seal ampule at room temperature. The ampule is attached via a glass to metal seal to a small CF flange. This flange is attached to the interaction chamber through two all-metal 1.5" UHV valves in series. The glass inner seal of the cesium ampule was broken during the initial preparation of the vacuum chamber, and it was at that time necessary to heat the ampule to about 100°C for a few days to produce enough cesium vapor to see the MOT. The two valves are in place so that if it should become necessary to replace or modify the cesium source, it can be done without breaking vacuum in the main interaction chamber. The valve closest to the main interaction chamber is always kept wide open. The other valve is usually kept loosely shut, i.e., nominally closed but not hand tight. Once every week or two, we let additional atoms into the interaction chamber by fully opening the outer valve for 10 minutes or so. This increases the background pressure of cesium in the main chamber substantially because of the high vapor pressure of cesium, $1.3 \cdot 10^{-6}$ torr at 25°C [Nesmeyanov63].

Our experiments are conducted with optical forces, and optical access to the chamber is necessarily important. The viewport flanges each have 1/4-20 tapped holes. Four 1/2" diameter rods attached to all of the larger windows but the one below the chamber support the wave plate and mirror mounts for the molasses beams. The interaction beam's retroreflecting mirror is also attached directly to its viewport flange.

2.4.2 Magnetic field control

Naturally, operation of a magneto-optic trap requires precise magnetic field control in addition to the optical fields that we have already described. The most important task is to create the strong field gradients that add the trapping process to the damping that results from optical molasses alone. The trapping fields are generated by a pair of anti-Helmholtz coils that surround the interaction region. A secondary task is to null out residual magnetic fields in the interaction region, which is performed with three orthogonal sets of Helmholtz coils. Each of these sets of coils is centered about the interaction region in the center of the main interaction chamber.

Helmholtz coils are pairs of identical current-carrying coils that have current flowing in the same direction. Usually, they are separated in position along their axis of symmetry by a distance that is equal to their radius. This configuration leads to an extremely uniform magnetic field over a large region. In practice, good uniformity of the magnetic field is obtained even if the Helmholtz condition (that the separation is equal to the radius) is not precisely met, or if (for example) square coils are used instead of circular. So-called “anti-Helmholtz” coils are Helmholtz pairs except that the currents in the two coils run in *opposite* directions.

The anti-Helmholtz pair is mounted directly to the interaction chamber, by Lucite clamps that are wrapped around the chamber flanges where the viewports for the vertical MOT beams are mounted. The coils themselves are wound around a water-cooled anodized aluminum frame. The coils are 6.2” in diameter and are wound with 202 turns of 24 gauge magnet wire. The coils are typically run at currents near 3 A, and produce field gradients of 11 G/cm along the (vertical) axis of the coils, and half that along the horizontal axes. The coils dissipate 80 W of heat into the cooling water, which flows at a rate of 10 gallons per minute through the frames.

The current through the coils is controlled by an interface box that drives the coils with push-pull pairs of high-current operational amplifiers. This configuration allows

independent control over the current in each coil, which can be used to shift the location of the magnetic-field zero. It is also capable of shutting off the 3 A of current in around 100 μ s, upon receipt of a control signal from the computer. In the course of our experiments it is important to be able to shut off the fields quickly when we shift from gathering atoms (in the MOT) to performing an experiment.

As it turns out, the magnetic fields decay over a much longer time scale than the current. There is a large amount of conductive material in the chamber, the flanges, and the aluminum coil frames, all of which contributes to eddy currents. The eddy currents lead to a relatively slow magnetic field $1/e$ decay time of order 3 ms. For most of our earlier experiments, the residual fields were simply not important because the far-detuned interactions are not affected strongly by the fields. When we later developed more sensitive methods of determining the magnetic fields, we learned that the residual magnetic fields in the interaction chamber are still changing in excess of 350 ms after the current in the coils is shut off. These fields are potentially very disruptive to our state preparation experiments, as we will see later. The source of this longer field decay timescale has not been positively identified, but we believe it is due to the presence of the ferromagnetic material Kovar, which is used in the glass to metal seals on our viewports.

The current controller box also controls the three pairs of Helmholtz coils that surround the chamber. These coils are in the shape of squares with truncated corners and are wound around anodized aluminum frames. The pairs are nested together and have sides of about 15" with separations of about 8". They are wound with 44 turns of the same wire that is in the anti-Helmholtz coils. Their field is estimated at 2 G/A, along their individual axes of symmetry. The primary purpose of these coils is to cancel out the ambient magnetic fields at the center of the interaction chamber. Besides the earth's magnetic field, there are also stray fields from our ion pump, from magnetic bases on the optical table, and from the optical table itself, which has a ferromagnetic surface. The resulting field is nearly vertical, with strength of about 0.8 G, and can be cancelled very well with the Helmholtz coils. The procedures for zeroing out the fields are not obvious. Although we

have a FW Bell Gaussmeter with Hall-effect probes, the probes cannot be inserted into the interaction region. We must instead rely on methods that use our atoms as probes. One method involves observing atoms drift in optical molasses. We gather atoms in the MOT and subsequently turn off (only) the magnetic fields. Without confinement, the atoms fall under gravity. The atomic distribution can be distorted or pushed by magnetic fields, since they can (for example) redefine the zero towards which atoms are attracted. This procedure is sufficient only to make sure that there are not gross remaining fields— it is not sufficient for experiments in which we are very sensitive to the fields. Direct measurement of the Zeeman splitting between the hyperfine sublevels is possible with methods that we discuss later.

2.5 Imaging

Aside from a few minor calibration procedures, all of the data that we take is acquired by analyzing pictures of atomic distributions. Our camera is model number TE/CCD-5122TK/1UV from by Princeton Instruments (now a division of Roper Scientific). The camera is a cooled, 16-bit monochrome CCD camera, with a resolution of 512×512 pixels. Internally, the camera uses a TEC to cool the temperature of the CCD chip to about -30°C , which substantially improves the noise characteristics. The cooling water that goes to the Ti:sapphire and the anti-Helmholtz coils also serves to remove heat from the hot side of the TEC. The camera is controlled by a massive Princeton ST-138 interface. The ST-138 powers the camera, controls its temperature, drives the shutter, digitizes the output signals, and provides a serial interface. The serial interface is read out by a NuBus card, which necessitates the presence in our experiment of a computer old enough to have a NuBus interface. (Modern interfaces are not available from Princeton.)

The shutter on the camera is mechanical, and is the only moving mechanical part on our table. The shutter can be triggered by a TTL signal sent to the camera controller along a BNC cable (that is, independent of the computer interface). With a flashlight and a photodiode, we measured the time for the shutter to fully open and close to be 5 ms and 12 ms, respectively. (The electronic signals must be delivered to the controller 15 ms

before the shutter must be fully open and 20 ms before it must begin closing.) In order to have a well-defined imaging time and uniform sensitivity, we only perform imaging when the camera shutter is fully open. The imaging occurs after the free-expansion time in our experiments when we freeze the atoms in optical molasses. In practice, we begin opening the shutter during the free-drift period, so that it is fully open when the imaging is to begin. We then pulse on the optical molasses light to expose the CCD chip in a well-controlled way. The shutter's mechanical nature has also been a source of concern simply because of the vibration that it produces. The shutter makes a sound that can be easily heard in our noisy laboratory when either opening or closing. The vibration that the shutter transmitted to the optical table was at one time strong enough to occasionally unlock the diode lasers. The repump laser, particularly because of its design, exhibits sensitivity to acoustic noise and vibrations. We constructed a mount for the camera that interleaves several layers of aluminum mounting structures with layers of Sorbothane (a viscoelastic damping material) that has significantly reduced vibrational coupling to the table. After a camera "click," there is still some signature of the vibration on the diode laser lock error signals, but the lasers are able to maintain their lock. Again, this is not a concern because the camera click is the very last thing in the experiment, and the laser lock signals are stable before the next experiment begins.

The camera points at the atoms through one of the small, horizontal windows on the main interaction chamber. The port of the interaction chamber directly opposite this window leads to the other vacuum components, but also to a final window that is in the line of sight of the camera. We covered this window with a piece of very black paper to provide a dark background for imaging atoms.

The camera has a standard Nikon-type lens mount, and we use a Nikon 105 mm f/2.8 D macro lens that reduces the image by 1:1.8 to fit on the CCD chip. We calibrate the distance scale of the camera by first focusing the camera on atoms trapped in the MOT. We then (without adjusting the focus or magnification) image a precision reticle distance scale, and replace the camera. With this procedure, we are able to directly measure the number

of pixels per millimeter.

Each individual experiment consists of a series of steps leading up to the time that we take a picture with the camera. After the picture is taken, we download the picture to a Power Macintosh 7100 computer that we programmed to interface with the ST-138 controller. The downloading process typically takes several seconds since the images are rather large and the controller (and its computer) is old. If it becomes possible at some point to load our MOT substantially faster, the camera readout time will become the rate determining factor in experimental repetition rate. The first step in processing the signal in the computer is to subtract a background signal. Typically we take 2 – 7 pictures of “the dark” with the CCD camera before each data run, several times per hour. These pictures are averaged together and the average is subtracted from every picture that we take in the data run after that. This eliminates a lot of the irregularities that the raw images display. The absolute gain of the camera also has a tendency to fluctuate over long periods of time, which is important only when we care about the total fluorescence count from the integrated atomic distribution. We have developed similar background-like methods to correct for these fluctuations when it is necessary.

The slow readout time of the Princeton CCD camera makes it impractical to use for everyday alignment procedures and monitoring of the MOT. For these purposes we have a supplemental CCD camera made by Ikegami that we keep trained on the MOT as well. The output of this camera is continuously displayed on a television monitor, and is one of the tools that we use most for optical alignment. At this point, let us note that the trapped atoms are, generally speaking, hard to see. Nearly all of the laser beams on the optical table are invisible, and when the MOT glows brightly, it does so in the near infrared. We do have several infrared viewers (Find-R-Scope, from FJW Optical Systems), but the trapped atoms are localized in a spatial region about the size of a grain of sand and it isn’t possible to get very close to the center of the interaction chamber. It is only with the use of our various cameras that we can really tell what is going on in the region near the trapped atoms.

2.6 Computer Control

2.6.1 Design philosophy

It is almost impossible to imagine conducting complex experiments of this scale without computer control of some sort. Certainly, many of the tasks could be controlled by an interconnected series of analog pulse generators, or possibly a very fast punched-card reader. It is our good fortune to work in an age where desktop computers can be programmed to rapidly execute complex sequences of events without human intervention. In our experiments, we have taken the attitude that anything that can be automated should be. The high degree of computer control in our setup means that most of the changes between two experiments can be made with programming changes only. A substantial part of my contribution to the experiment as a whole has been the design and programming of the data collection software.

There are several different ways to control an experiment, depending upon the specific requirements. One way is to control each event as it happens in real time. The computer can send out specific signals, in a specific order, take data from various sources, and make changes to the experiment based upon the data that it receives. This is a common method in many laboratories. It is not, however, practical for our experiments because the response time of the computer may depend upon many factors. Interactive control by a computer can lead to unpredictable and substantial delays, sometimes comparable to the timescales of our experiments. We require a high degree of precision in the timing and generation of many signals. Our solution has been to program a number of independently running instruments from the computer. Each of these instruments is fully programmed *before* the experimental sequence begins, and executes its program in step with a common external clock signal.

2.6.2 Implementation

There have essentially been two generations of the data acquisition (DAQ) software and electronics. Since they are not terribly different, we will describe both of them here. The experiment is controlled by custom data acquisition software written in LabVIEW from

National Instruments Corp. (NI). The LabVIEW environment has been an excellent tool for programming, and it kept up well with our changing needs, even as the programs have become much more complicated and extensive over time. The experiment is run on a Macintosh platform. Originally, we used the single Power Macintosh 7100 computer to both control the experiment and take data from the computer. In the later generation of the DAQ setup, we have added a faster Power Macintosh G4 computer that handles all aspects of the experiment except interfacing with the Princeton Instruments CCD camera. The two computers communicate and exchange data over our ethernet network. Some of the data transfer is handled directly with TCP/IP routines that are built into LabVIEW.

Most of the external instrumentation that controls the experiment is programmed directly or indirectly over GPIB (general purpose interface bus; IEEE-488). Each data taking computer has an NI GPIB controller card: NB-DMA-2800 on the 7100, and PCI-GPIB on the G4. When the newer computer was installed, all of the instrument programming tasks were transferred to it. At the present time the GPIB chain includes the two computers, nine arbitrary waveform generators, two digital delay generators, an oscilloscope, two counters, and several other instruments. The timebases of these instruments, including the waveform and delay generators, are synchronized to a 10 MHz reference signal. The reference signal comes from an EFRATOM model LPRO rubidium (atomic) clock that is further distributed by a Stanford Research Systems (SRS) model FS710, 7-channel 10 MHz distribution amplifier.

The arbitrary waveform generator (AWG) is one of our most powerful tools. In many circumstances, we need to provide complicated waveforms as a function of time to execute the experiment. Any waveform that needs sufficient speed of execution or cannot be expressed as a pure binary on-off switching operation has been handled by these machines. A typical waveform might be a series of pulses with different amplitudes for a noisy kicked-rotor experiment. We can generate the series of pulses on the computer and download an array of amplitude versus time values to one of the AWG units. Upon receiving a digital trigger, the AWG unit begins to output its buffer. The buffer executes in lockstep with

all of the other instruments because of the precision timebase. In the kicked-rotor example that we are using, the output signal goes to the input of an IntraAction AOM driver that controls the interaction beam intensity.

Eight of our AWG units are model DS345 from SRS. This particular model stores up to 16,299 points in its buffer, which it can execute at up to 40 MS/s (megasamples per second). The amplitude of each output point is in the range $\pm 5\text{V}$, determined with 12-bit resolution. We have used these units in many applications, such as controlling the interaction beam intensity, the MOT beam intensity, and so on. Our other AWG is a newer unit from Agilent, model 33250A. This is a competing product to the SRS boxes, and it can store up to 65,535 data points and output them at up to 200 MS/s. We have used this synthesizer to produce the longest and most complex waveforms that we have used for state-preparation and chaos-assisted tunneling experiments.

Besides the AWG units, we also employ two model DG535 digital delay generators from SRS. These units can each output two fully independent pulses defined by four logical edges. Because each of the output lines can be set independently to an arbitrary analog level, we have employed them as level generators for calibrations in addition to their role as pulse generators.

One of the DG535 units serves as the main triggering device of the experiment as a whole. Typically, we begin the experiment with a trigger issued by the data taking program, called the *software trigger*. After this trigger, all events are handled by external timing. The “trigger” DG535 then issues a series of triggers to the other instruments. One trigger indicates the beginning of post-MOT cooling procedures, and another initiates the beginning of the interaction sequence. A set of multiple triggers is important because the buffers of the AWG units can be exhausted quickly when running at high speed. When we run the DS345 at 10 MS/s (100 ns data point spacing), the buffer is exhausted in 1.63 ms. It is obviously necessary to trigger that instrument exactly at the moment that we wish the interaction to begin, so that it does not waste precious space in its buffer.

Besides the GPIB instruments, there is one additional class of instruments that we have used to control the experiments. Both data-taking computers have DAQ interface boards from NI. The 7100 has a model NB-MIO-16L-9 I/O board. This board incorporates analog and digital I/O capabilities, including analog waveform generation. It also includes digital counters that can be used as triggers or configured as a multichannel digital pulse generator. While the capabilities of this board are impressive and potentially very useful, it does have a fatal flaw. The various sections of the interface cannot be synchronized or simultaneously triggered. In our design philosophy where we *do not tolerate* timing errors, this renders the board approximately useless. For the set of experiments that was based solely on the 7100, we did end up using up to three of the counters as pulse generators. We used them to issue the first trigger of the experiment, as well as for digital timing that was not extremely time critical, such as switching off the anti-Helmholtz coils at the appropriate time in the experiment. In more recent times, we have used this board primarily for the analog input channels, which are the only ones available in our experiment. These channels are used to check for Ti:sapphire laser power drifts during long data runs, and in many types of calibration procedures.

In the later experiments, we have relied heavily upon a PCI-DIO-32HS digital interface from NI that resides in the Power Macintosh G4. This board has 32 channels that can be configured for input or output in a variety of modes. A particularly nice feature of this board is that it can operate in *pattern generation* mode. In this mode, the board outputs one word of data, up to four bytes, at a time. Upon receiving the proper signal, it can then output the next word in its buffer. The board is quite fast, and in principle, the total length of the buffer is only limited by the computer's memory. Much like one of the external instruments, the entire buffer for this board is written well before the experiment begins. The buffer is triggered during the experiment by an input signal from the DG535, and each new data point is output in step with an external clock. In testing the speed of the various interface components, we found that the interface did not always succeed in outputting a new data point if the sampling rate was too high. The board works reliably if

the input clock rate is less than or about 5 MHz. As it turns out, for most of the cases where we need digital control, we do not need this much speed. We usually clock the data out at a sample rate of 1.25 MHz. The output clock is provided by an external box that counts down the 10 MHz reference clock signal (with a ripple counter chip) to one of several selectable lower frequencies and outputs a TTL signal. The board itself has a high-density computer connector that does not suit itself to experimental interfacing. We built an interface box for the board that buffers all 32 output channels with 50 Ω output drivers and standard BNC connectors. (A similar interface box provides buffering and connectors for the NB-MIO board.) Currently, we are using 16 output channels of this board. The signals that this board creates have proven surprisingly versatile. We initially used it to produce those signals that had been handled by the NB-MIO board. Some of the channels have been used to “gate” the RF switches that we keep in the AOM driver boxes. Recently, we have also used these signals to control a bank of ten RF switches to switch various electronic signals in the experiment. We have used these, for example, to control AOM drivers by switching the input between ground and an analog output from one of the DG535 units. This sort of versatile control has proven invaluable in controlling more complex experimental sequences.

The data taking program itself has evolved considerably. In the early experiments, up through the time that we added the second computer, we had a separate program that controlled each experiment. The programs have typically consisted of giant layers of nested loops that control the instruments and collect the data. As an example, we had one program to perform temperature measurements, and a separate program to perform kicked-rotor type experiments with amplitude noise. There is a large amount of common ground in all of the experiments that we conduct—most of the experiment is the same, no matter what physics we are currently studying. With the large number of independent code bases that I was maintaining, it eventually became extremely time consuming to fix every bug that we found in all the different programs. There were several other problems with the initial design of these programs as well. The basic user interface that these routines had was a “front panel” like that of a laboratory instrument. The panel had places to input each of the

parameters for the data run and would display relevant data plots and output parameters as they became available. This system was functional, but very difficult to change. Adding any single new variable or parameter to the system was an extensive chore and required altering the visible interface of the program. This type of problem led to the creation of a lot of “hidden variables,” parameters that were only visible in the code and not in the output of the program. At the time that we added the second computer and the digital interface, I took the opportunity to write the data taking program over from scratch.

Perhaps most importantly, the new program has a consistent set of interfaces. Each possible experiment type is chosen from a single pop-up menu, and more can be added. The differences between all of our experiments are quite minor in the scope of the program. The various parameter sets for the different experiment types are also chosen from pop-up panels, and it is relatively easy to manage the several hundred parameters that govern an experiment in its entirety. There are still some hidden variables, and they do occasionally cause trouble, but their number has been reduced. Data is output on a general interface that is also common to all of the experiment types. Most of the analysis that we perform on the atomic distributions is performed at the time of data collection, and the data is saved to disk for later consideration.

2.7 An Experiment in More Detail

A timing diagram for a sample kicked-rotor experiment is shown in Fig. 2.12. This diagram shows the typical signal assignments and sequence of events for our experiments up through the summer of 1999, when we added the second computer to the DAQ system. The essential sequence of events is that which we described in the overview (§2.1), and is indeed a fairly general picture of our procedures.

The sequence actually begins before the diagram shown here, with the programming of the GPIB instruments and the I/O board in the computer. Typically the full programming sequence takes less than a second. For some of the more complex experiments that we will describe later, the programming can take much longer. For example, to program the entire

Timing Diagram (5/99) -- Cs δ -Kicked Rotor




Figure 2.12: Timing diagram for a kicked-rotor experiment. This shows the signals that control a typical experiment. For each instrument, the trigger source is indicated, and the time at which the trigger occurs is labeled by a small “T.” The DS345 arbitrary waveform generators each have one output, and the DG535 delay generators have four outputs labeled by the letters A,B,C,D. The time intervals shown here are not to scale and represent very different time scales. For example, the loading period is about 6 s, and the pumping time is about 50 μ s. A timing diagram for the more complex experiments involving quantum state preparation is shown in Fig. 4.9

waveform memory of the Agilent 33250A takes roughly 18 s. We have, however, developed a number of shortcuts to speed up the programming process. For example, rather than programming every instrument before each shot of the experiment, they are all programmed once at the beginning of a data run. Then, before subsequent shots we only reprogram those

instruments that absolutely need to be. In special circumstances, we can use more complex tricks to avoid long reprogramming stages. A particularly effective trick can be used when programming the Agilent 33250A for an experiment in which we sweep the total interaction time. Instead of programming the exact sequence to be executed before every shot, we instead program the Agilent box once, with the full sequence for the maximum interaction time that we are interested in. We then, before every shot of the experiment, simply step the time (with a digital line) at which an RF switch disconnects the signal that otherwise runs from the Agilent box to the AOM controller. Tricks like these are necessary because the effective signal to noise ratio of our experiments is determined by the number of times that we can execute an experiment. It is almost always worth working hard to improve the repetition rate of the experiment.

When the instruments are programmed, they are initialized to the values that begin the experiment. Among other things, the interaction beams are off at this time and the signals are set so as to load the MOT. The optical molasses and repump beams are on, at their full intensities. The detuning of the optical molasses beams is at a value that is optimal for loading a large number of atoms into the trap, typically about -15 MHz. The anti-Helmholtz coils are on, and the camera shutter is shut. As long as we have these ingredients, the MOT is operational and we are gathering atoms from the background vapor. Often it is advantageous to gather as many atoms as we can from the background vapor, and it is fine to begin loading the MOT before the experiment is “officially” triggered. In cases where we are concerned about the absolute fluorescence from the atoms after some interaction, this is not acceptable, and we do not turn on the magnetic fields until a specific trigger is issued.

The experiments are initiated by the software trigger, which is issued by the DAQ software to the relevant timing hardware. In the case of this setup, that hardware is counter 5 on the I/O board. After a delay of 5 ms, the counter board sends the “main trigger” signal which triggers the DG535 delay generator at GPIB address 11. Thereafter in the experiment, this DG535 controls the timing. A secondary delay period of 5 ms begins when

the main trigger is issued, and the “official” MOT loading time begins. The MOT is loaded for 5 – 7 s, during which the MOT approaches the long-time limit of atom number, which we estimate to be of order 10^7 atoms.

After the loading phase has ended, we issue the *cooling trigger* and begin the cooling phase. In these simple experiments, the cooling phase consists of a 3 ms polarization gradient cooling interval where we detune the optical molasses light further from resonance to between 35 and 55 MHz, and decrease its intensity to (typically) 40% – 60% of that which we use to load the MOT. The particular values have often changed depending on our particular requirements. This procedure dramatically reduces the final temperature of our atoms. Once this cooling phase has ended, the optical molasses light is shut off and the current through the anti-Helmholtz coils is shut off.

Generally, it is in the period after the cooling trigger that we perform the state preparation as well. In the simple case that we are discussing here, the only additional state preparation is a brief stage of optical pumping. In order to ensure that all of the atoms are in the $F_g = 4$ ground state, we leave the repump light on slightly longer than the optical molasses light. An optical pumping period of 50 μs is sufficient to promote the few remaining $F_g = 3$ atoms to $F_g = 4$.

When we have concluded these procedures, we have an atomic sample that is ready for an experiment. The momentum spread is near Gaussian with a typical width of $\sigma_p/2\hbar k_L = 4$, or about 12 μK . The distribution in position is also nearly Gaussian, with a width of $\sigma_x = 0.15$ mm. The width in both p and x is roughly the same in all three directions, however we are usually only concerned with these values as measured along the axis of the interaction beam. The *interaction trigger* is then issued, which typically triggers an arbitrary waveform generator. In our timing diagram, a DS345 produces a series of short pulses that last for up to 1.6 ms, a value that is limited by the buffer length of the DS345. The atoms interact with the one-dimensional standing wave, gaining a characteristic momentum distribution that we must now detect. During the interaction period, we “gate” the interaction beam control signal on with an RF switch. A second RF switch is used to

ensure that the repump beam is fully off during the interaction time.

A free drift period of 15 – 40 ms follows the interaction, and it is the ballistic expansion during this time that allows us to determine the momentum distribution. During many experiments we sweep the total interaction time. For example, we might step through the number of kicks in a kicked-rotor experiments with interactions that last between 0 and 1.6 ms. In cases like this, there is some degree of ambiguity in the exact definition of the free drift time. Because of this, we always define the free drift to begin at the *beginning* of the interaction time. The size of the freezing molasses region that we image and the length of our drift times limit the total momentum scale of our measurements to about $\pm 80 \cdot 2\hbar k_L$.

At the end of the specified free-drift period, we turn on the freezing molasses to expose the CCD camera and measure the atomic distribution. The camera shutter is triggered during the free-drift period so that it is fully open by the time that the exposure begins. The typical exposure time is 10 ms, although we have occasionally used longer exposures. After the end of the exposure, we begin a reset period. The reset delay allows time for the camera shutter to close before light is reintroduced to the chamber. After this delay, the MOT is reset for the next experiment.

After the experiment is over, the data is read out from the CCD camera and integrated along the vertical axis to form a one-dimensional momentum distribution along the direction of the standing wave. Finally, the data is plotted and the computer prepares the instruments for the next experiment.

2.8 Temperature Measurements

One of the most important standard calibrations that we perform is the temperature measurement. This is a standard diagnostic tool that we use to make sure that all of our experimental systems are functioning properly. Furthermore, we use the temperature as a standard number by which we evaluate the effectiveness of cooling techniques and our understanding of the initial conditions for our experiments.

The momentum distributions that our MOT produces are described fairly well by a classical thermal model. If we assume a Maxwell-Boltzmann distribution, the probability P of an individual atom having kinetic energy $E = p^2/2m$ is

$$P(E) = \frac{e^{-p^2/2mk_{\text{B}}T}}{\int_{-\infty}^{\infty} dp' e^{-p'^2/2mk_{\text{B}}T}} = \sqrt{\frac{m}{2\pi k_{\text{B}}T}} e^{-p^2/2mk_{\text{B}}T} = \mathcal{A} e^{-p^2/2\sigma_p^2}, \quad (2.1)$$

where k_{B} is the Boltzmann constant and m is the atomic mass. We have defined the momentum variance as $\sigma_p^2 \equiv mk_{\text{B}}T$, and \mathcal{A} is effectively a normalization factor. The procedure for the measurement is essentially the sequence that we have previously described, when the interaction time is set to zero. The momentum distribution is determined from the final position distribution and the drift time, assuming that the initial position distribution is very small compared to the final distribution. When the distribution is very cold or the drift time is short enough, this may not be a good assumption. Since the position and momentum distributions are both (nearly) Gaussian, the total distribution is given by their convolution which is also Gaussian. The total width in position after a drift time t is

$$\sigma_x(t) = \sqrt{(\sigma_{x_0})^2 + (\sigma_p t/m)^2}, \quad (2.2)$$

where σ_{x_0} is the size of the initial position distribution. In a typical temperature measurement, we measure the position distribution at eight different drift times between 15 and 40 ms. We fit the values $\sigma_x(t)$ to the model above to determine σ_{x_0} and σ_p . We can then determine the temperature by using the definition of σ_p . It is important to note that the distribution is not *entirely* thermal. The simple Gaussian distribution that we fit to the distribution only accounts for about 96% of the atoms. Besides this, there is a weak pedestal with a roughly exponential distribution. This is a side effect of polarization gradient cooling, and has been observed in other groups that use a similar setup.

A sample set of data from a temperature measurement is the 10 ms exposure time curve from Fig. 2.3. Note that in this data, the curve (for the short detection times that we use) is nearly linear, suggesting that the initial spatial distribution does not make a major contribution to the final distribution. For other types of measurements, we do not deconvolve the initial spatial distribution.

Often, we simultaneously measure the temperature in both the vertical and horizontal directions. A secondary piece of information that is yielded by the temperature measurement is the vertical position of the atomic cloud as a function of time. If this data matches the curve that we expect for gravitational free fall, it confirms that the atoms are indeed under free fall, i.e., not launched by optical or magnetic forces. Furthermore, it provides an independent verification of the spatial calibration of the imaging system.

Chapter 3

Experiments with the quantum kicked rotor

3.1 Overview

We have already introduced many of the physical phenomena that occur in the kicked rotor system. We have also introduced our experimental setup and procedures. With this background in place, we will now describe our first major set of experiments. First we will address some differences between our experiment and the ideal δ -kicked rotor. The most significant difference is the nonzero length of our pulses, and it turns out that the short pulses in our experiment model δ -functions very well within a limited range of momentum. We have used our kicked rotor system to observe dynamical localization, quantum resonances, and the effects of short-term correlations on quantum dynamics. We have already seen that there are vast differences between the behavior of quantum and classical kicked rotors and so this is an ideal system in which to study quantum-classical correspondence. The quantum coherences that create interference effects can be destroyed by noise or dissipation, which we can introduce to the experiment in a controlled manner. Finally, we demonstrate that a noise-driven quantum rotor can be driven hard enough that its behavior is indistinguishable from classical behavior.

3.2 Finite Pulses and Dynamical Localization

When investigating the dynamics of the kicked rotor in Chapter 1, we specialized to kicks that were described by Dirac δ -functions. This simplifies the equations of motion substantially and allows us to derive simple mappings that describe the behavior. The pulse duration must be greater than zero in any real experiment.

In most of our kicked-rotor experiments, the pulses are nearly square, with a pulse width of $t_p = 300$ ns. There are several reasons why we cannot make the pulses much shorter than this. The most fundamental reason is that we are limited by the total amount of laser power that we have. To maintain constant interaction strength, we must increase the intensity of the interaction beam when we make the pulses shorter. A second problem derives from the resolution of our arbitrary waveform generator and AOM system. The sequence of pulses is controlled by an AWG such as the SRS DS345. For a typical experiment of 80 kicks separated by $T = 20$ μ s, the waveform memory is exhausted when the sample rate is 10 MS/s, or one point every 100 ns. An individual kick is then described by only three points in the waveform memory, and it would be difficult to make the kicks much shorter. It is not absolutely necessary that we use an AWG to generate the kick sequence, however it does have many technical advantages over other methods that we have considered. Even if we were able to generate much shorter *electronic* pulses, the response time of the interaction beam AOM is still of order 20 ns. A final reason is that in the limit of extremely short pulses, the adiabatic assumption that we invoked in the derivation of the potential is no longer valid.

The fundamental difference between our experiment and a true δ -kicked rotor is that an atom can move a finite distance during a finite-duration kick. If an atom is moving fast enough, we can imagine that it would travel over an entire period of the potential during the kick. If this happens, the net force on the atom will average to zero. An atom at this speed does not feel any kick from the standing wave, and the effective stochasticity parameter K_{eff} goes to zero. We then have a momentum-dependent stochasticity parameter, and it is reasonable to expect that it is only below a certain total momentum “boundary” that the system approximates a δ -kicked rotor. A simple estimate of the boundary, p_b , comes from the condition of the distance travelled during the kick being equal to the period of the potential. We then have (in physical units) $p_b t_p / m = \lambda / 2$, or

$$p_b / 2\hbar k_L = m\lambda^2 / 8\pi\hbar t_p. \quad (3.1)$$

This simple *classical* model of the boundary tells us that the boundary is much higher for

cesium atoms than for sodium atoms (with the same pulse width), because of the higher mass and transition wavelength of the cesium atom. The location of our boundary is a tremendous advantage over the early kicked rotor experiments performed with sodium atoms, and is the primary reason that the cesium experiment was constructed.

We have seen, from a classical argument, that there is a certain momentum at which diffusion simply shuts off. The momentum boundary is a purely classical effect that arises from the dynamics of the system. One reason that it is important to understand this effect is that we wish to study dynamical localization, which is also an effect that suppresses momentum diffusion. We must be very careful to distinguish between classical effects and dynamical localization.

Let us now consider the finite pulse effects in more detail. We begin with the δ -kicked rotor Hamiltonian in scaled units, $\mathcal{H} = \frac{p^2}{2} + K \cos x \sum_{n=-\infty}^{\infty} \delta(t - n)$. We can rewrite this as the discrete Fourier series

$$\mathcal{H} = \frac{p^2}{2} + K \sum_{n=-\infty}^{\infty} \cos(x - 2\pi nt), \quad (3.2)$$

so that the potential is equivalent to an infinite number of constant-amplitude cosine terms moving at different velocities [Reichl92]. Each cosine term leads to a primary nonlinear resonance at velocity $p = \dot{x} = 2\pi n$. From this expansion we can see that the δ -kicked rotor system is 2π periodic not just in position, but in momentum as well. For finite-length pulses, the Fourier transform is no longer exact and the effective value of K at each resonance m depends upon the velocity. If we approximate our pulses to be square, the Fourier transform of the pulse shape is a sinc function ($\text{sinc}(x) \equiv \sin(x)/x$), and

$$K_{\text{eff}} = K \frac{\sin(pt_p/2)}{(pt_p/2)}, \quad (3.3)$$

where t_p is the scaled pulse width. In physical (non-scaled) units, this factor is the duty cycle $\alpha_{\text{pulse}} \equiv t_p/T$. The actual shape of our pulses is not particularly square— it is effectively a low-passed square pulse function (see Fig. 3.12). However, it turns out that the experiment is not highly sensitive to the pulse shape, and this model is usually sufficient. In cases




Figure 3.1: Calculation of the effective stochasticity parameter, K_{eff} , as given by Eq. 3.3, for several pulse widths. For each case, $K_{\text{eff}}(p = 0) = 10.5$, which means that the well depth has been lowered for the longer pulses. The horizontal line is the result for true δ -function kicks. Below this in successive order, are finite pulses with scaled pulse width (duty cycle) 0.014, 0.024, 0.049, and 0.099. The 0.014 case corresponds most closely to typical experimental parameters. The curves are plotted only until their first zero.

where the exact pulse shape *is* important, we rely upon numerical simulations that directly incorporate the pulse shape.

The model that we have described shows that the effective stochasticity parameter gradually drops off with increasing momentum, when the pulse width is constant (Fig. 3.1). If we wish to define a hard boundary in momentum, a natural choice is when K_{eff} approaches unity. Above this value, global diffusion is possible, and below this value stable structures in the phase space dominate the behavior. For the values of $t_p = 300$ ns and $T = 20$ μ s, the




Figure 3.2: The momentum boundary in phase space. Classical surfaces of section for the kicked rotor are presented for three kick widths, all with $K = 10.5$. One unit cell is shown in position, and a wide range is shown in momentum. In the δ -kicked rotor (a), the periodicity in momentum is exact, and the momentum scale spans many unit cells in momentum. For pulses of finite duration, the momentum periodicity is broken. For a scaled pulse width (duty cycle) 0.014 (b), the dynamics resemble those of the δ -function case within a bounded region. This value of the pulse width is typical of our experiments, and the boundary occurs well outside the region $\pm 80 \cdot 2\hbar k_L$ that we use in our experiments. For a longer pulse width with duty cycle 0.049 (c), the boundary is much closer to zero. Note that there are regions beyond the main boundary for which diffusive behavior is expected. These regions occur where $K_{\text{eff}} > 1$, and fill out a pattern suggestive of the sinc function (Eq. 3.2).

duty cycle is 0.015, and the boundary occurs near $\pm 200 \cdot 2\hbar k_L$. Note that this is much larger than the highest momentum that we can detect, which is $\pm 80 \cdot 2\hbar k_L$. Nonetheless, there is an appreciable falloff in K_{eff} , which approaches a 25% reduction within the range that we can image. We plot the phase space for the kicked system with pulses of three different durations in Fig. 3.2.

We now come to our first experimental results. By changing the pulse width we




Figure 3.3: Observation of dynamical localization and the classical momentum boundary. Atomic momentum distributions are shown for the kicked rotor after 0 kicks (heavy solid line), 17 kicks (dotted), 34 kicks (dashed), 51 kicks (dash-dot) and 62 kicks (solid). For cases (a-c) the well depth was scaled to keep the stochasticity parameter K nearly constant. The intensity scale is logarithmic and is in arbitrary units. For short pulses in (a) and (b), the system exhibits dynamical localization. For the other two cases, the behavior is dominated by the effects of the classical boundary. Case (a) is typical of our experiments. The indicated value of α_{pulse} is the duty cycle for each case.

can vary the location of the boundary. When the pulse width is short enough, our system is an excellent approximation of the δ -kicked rotor within a limited range of momentum, and we can observe dynamical localization. The momentum distributions from one of our experiments are shown in Fig. 3.3. Each curve in this plot represents an independent experiment, and is the integrated, one-dimensional momentum distribution as determined from a 15 ms free-expansion period, without deconvolving the initial spatial distribution. The kicks were spaced apart by $T = 20 \mu\text{s}$, and the detuning of the interaction beam was 6.12 GHz to the red of the cycling transition. (This detuning was used for all of the experiments described in this chapter.) Our initial condition for the experiments is shown as the result after 0 kicks. This is a typical profile of atoms cooled in the MOT. About 96% of the atoms fit into a Gaussian with width $\sigma_p/2\hbar k_L = 4.4$, however there is also a broad, dim pedestal that has an exponential profile. A more detailed description of this distribution is available elsewhere [Steck00].

We now follow the time evolution of the system for several different kick widths. Our first kick width is near 300 ns and we are able to observe dynamical localization. During the kick sequence the momentum distribution rapidly changes from Gaussian to exponential. Once it has done so, the distribution essentially stops spreading. There is some residual energy growth, however it is minor compared with the initial growth. For slightly longer pulses, the dynamics are not affected in any visible way. When the location of the boundary enters the region of momentum that we can detect, its presence is obvious. Diffusion occurs out to higher momentum values where the dynamics suddenly become stable (near $K_{\text{eff}} = 1$) and diffusion halts. The boundary has a very different signature than dynamical localization, a sharp cutoff, rather than exponential decay in diffusion. Additional details about the classical boundary in our experiment are presented in [Klappauf99].

One word of caution is in order about the distributions displayed in Fig. 3.3. We did not employ an active optical pumping stage in our experiments at the time that this data was taken, and a number (that could possibly be as high as 20%) of the atoms were actually in the $F_g = 3$ ground state at the beginning of the experiment. These atoms are

effectively detuned 9.2 GHz further from resonance because of the $F_g = 3, 4$ ground state splitting, and did not interact as strongly with the standing wave. These atoms tend to remain near $p = 0$ and form a small bump that is especially visible for the localized cases in the long-time limit. We did not identify this problem until after the publication of one of our other papers, and resulted in an erratum [Klappauf98b]. A “cleaner” example of an exponentially localized momentum distribution is visible in Fig. 3.5. The $F_g = 3$ ground state population is also present in two data sets of energy evolution curves (Figs. 3.4, 3.9). The mixed population leads to a systematic reduction in the energy of order 20%, but does not otherwise affect the curves.

3.3 Effects of Correlations

In §1.2.3 we discussed how short-term correlations influence the overall diffusion properties of the classical kicked-rotor system. The presence of quasiaccelerator modes, streaming behavior and other correlation effects lead to oscillations in the diffusion rate as a function of K .

Our system is quantum mechanical and so we expect it to exhibit diffusive motion only for short times, after which the system settles into an exponential distribution that is characteristic of dynamical localization. It is then reasonable to expect that the characteristic width (the *localization length*) of the localized distribution depends upon the initial diffusion rate. The short-time quantum diffusion rate D_0 has been predicted by Shepelyansky to follow the classical predictions (Eq. 1.15) for $D(K)$ when K is replaced by [Shepelyansky87; Cohen91]

$$K_q \equiv K \frac{2}{k} \sin \frac{k}{2}. \quad (3.4)$$

The value K_q represents a *quantum* stochasticity parameter, and the short term quantum diffusion rate is then approximately

$$D_q(K, k) = \frac{K^2}{2} \left(\frac{1}{2} - J_2(K_q) - J_1^2(K_q) + J_2^2(K_q) + J_3^2(K_q) \right). \quad (3.5)$$

This formula then tells us that the classical oscillations in $D(K)$ appear in the quantum

system, but their location is scaled by the value of \tilde{k} . Zeros of Eq. 3.4 occur when $\tilde{k} = 2\pi j$, for integer j , which we recognize as the condition for the quantum resonance (§1.3.3).

We do not have a simple model for the momentum distributions that are expected in a regime of classical anomalous diffusion, however there has been some theoretical work in this area. Hanson, Ott, and Antonsen have suggested that in certain regimes tunneling out of the accelerator modes can be an important process in this system [Hanson84]. Beyond this, it is expected that anomalous diffusion in the classical dynamics leads to fluctuations in the localization length [Sundaram99].

We can study quantum diffusion rates in our experiments by simply stepping the value of K . One of the best measures that we have for quantifying the diffusion rate is the distribution energy $E = \langle (p/2\hbar k_L)^2 \rangle / 2$, which is obtained by directly integrating the momentum distribution. The energy figure reflects changes in the localization length and generally quantifies the amount of momentum that has been transferred to the atomic ensemble. We have not found it to be practical to fit the distributions directly to an exponential lineshape, and it is particularly inappropriate in some of our cases, where the distribution is not purely exponential. While the energy measurement is convenient, there are some problems with it. The energy is particularly sensitive to noise in the distributions because $\langle p^2 \rangle$ is most sensitive to the behavior at high momentum.

We trace out the $D_q(K)$ curve by measuring the energy (after a constant number of kicks) for kick sequences with different intensities. In order to confirm the quantum scaling law (Eq. 3.4), we can repeat this measurement at different values of $\tilde{k} = 8\omega_r T$ by changing the period between pulses. In this series of experiments, we swept the pulse period T from 10 to 60 μs , which corresponds to \tilde{k} in the range of about 1 to 2π . We chose a sequence of 35 pulses. The pulses themselves were measured to have full width at half maximum (FWHM) of 283 ns, with less than 3 ns variation. The typical rise/fall time of the light intensity is near 75 ns, and this slow timescale is responsible for the effective width of our pulses, which are nominally generated as 300 ns pulses.

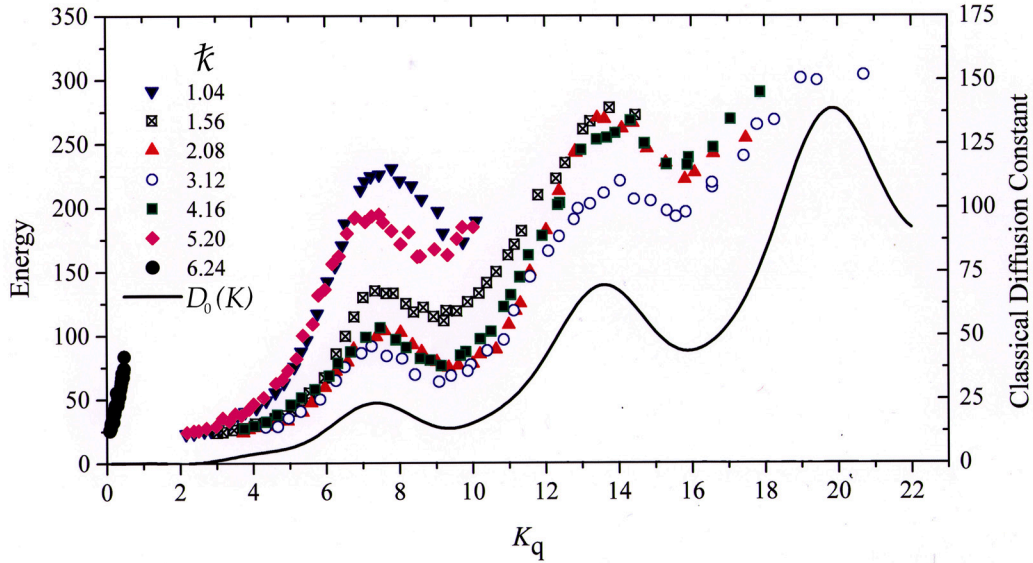


Figure 3.4: Quantum diffusion curves at various $\hbar k$. The measured energy $\langle (p/2\hbar k_L)^2 \rangle / 2$ is plotted as a function of the scaled stochasticity parameter K_q for seven values of $\hbar k$. We have also plotted the Rechester and White expression for the classical diffusion constant $D(K)$ (vertical scale on the right). All but one of the experimental curves match up very well with the classical expression after rescaling. The other curve, for $\hbar k \approx 2\pi$ is quite different because it is near the quantum resonance.

The results from our experiment are presented in Fig. 3.4. Here we see that the quantum diffusion rates match up with the Rechester and White expression for $D(K)$ (Eq. 1.15), which serves to confirm Shepelyansky's scaling factor (Eq. 3.4). Let us briefly discuss some of the systematic uncertainties in these measurements. The most difficult parameter for us to calibrate in any experiment is the absolute intensity of the interaction beam (see §2.3.3), which gives us an uncertainty of $\pm 10\%$ in K . As we have also seen, finite-pulse effects lead to an effective reduction in K that is of order 25% at the edges of our detection region ($\pm 80 \cdot 2\hbar k_L$). Beyond these effects, there is also uncertainty in the detection from the initial spatial extent of the distribution, from the spatial calibration of the CCD camera system, and from the ambiguity in the definition of the free drift time. Additionally for cases with high momentum (either in the regions of classical anomalous diffusion or where $K_q > 15$) the distributions tend to reach the edges of our detection region before the end of the kicking sequence. These factors, combined with the sensitivity that the energy measure




Figure 3.5: Effects of correlations on momentum distribution evolution. These two cases show the dynamics with $k = 2.08$ at the first peak and valley of the diffusion curve shown in Fig. 3.4. The first case (a) was measured with $K_q = 9.1$, and corresponds to the first true minimum of the diffusion curve. At this case, we observe exponential distributions that confirm dynamical localization. The second measurement (b) was performed with $K_q = 7.9$, near the peak of the diffusion curve. The distributions here are broader, and the distribution is no longer quite as simple. The time steps shown are 0 kicks (light solid), 5 kicks (bold), 10 kicks (dashed), 20 kicks, 45 kicks, and 70 kicks (heavy bold). The vertical scales are logarithmic and are in arbitrary units.

has to high momenta may lead to systematic errors on the order of (20-30)%. Besides these general errors in determining energies, the data shown in this plot have the residual $F_g = 3$ ground state population, which reduces the overall energy values but does not affect the locations of the peaks.

Beyond simply looking at the energy values, we have also examined the momentum distributions at the first maximum and subsequent minimum of the curve shown in Fig. 3.4. The momentum distributions at these values of K_q are compared in Fig. 3.5. At the minimum of the diffusion curve $D(K)$, we observe that the system evolves to exhibit dynamical localization. At the maximum, in the regime of classical anomalous diffusion, the behavior

is qualitatively different. The distribution is no longer described by a single exponent, and indeed it is not obvious that we should expect the system to display exponential localization.

To our knowledge, this experiment represents the first observation of the oscillations in diffusion rate in our system. Furthermore, it is the first confirmation of the scaling factor for the quantum kicking strength K_q . The curves in Fig. 3.4 have a well-defined dependence on K , and in subsequent experiments we have used this feature to directly calibrate our well depth. Additional details about this experiment are available in [Klappauf98a].

3.4 Quantum Resonances

We have preliminary experimental evidence in Fig. 3.4 that the behavior of the quantum kicked rotor can change dramatically at certain values of $\hbar k$. In particular, the quantum short-term diffusion rate scales along with the classical diffusion rate, except when $\hbar k$ is near 2π . As we have seen earlier (§1.3.3), a quantum resonance occurs when this condition is met. This effect is in some sense more profound than dynamical localization— at quantum resonances the entire interaction collapses to a single effective kick or averages to zero at the antiresonance. When this is the case, the dynamics that would otherwise generate localization are no longer present.

In our earlier discussion of quantum resonances from a theoretical standpoint, we noted that the initial distribution may have a strong influence over the types of behavior that we will see. Fig. 3.6 compares the behavior at the quantum antiresonance ($\hbar k = 2\pi$) with two different initial conditions. Here, the free evolution factor U can be 1, -1 , or something far less interesting, depending on the initial momentum state. The plane wave initial condition ($p_0 = 0$) shows periodic recurrences where the phase factor flips sign between successive kicks. We also show the time evolution of a system that begins with a continuous momentum distribution close to that of our experiment. This case evolves very differently. There is a nearly static central region flanked by two weak ballistic peaks. It is clear from these results that the choice of initial conditions can make a large difference in the observed dynamics. Indeed, the primary signature of quantum resonances in the earlier sodium experiments was

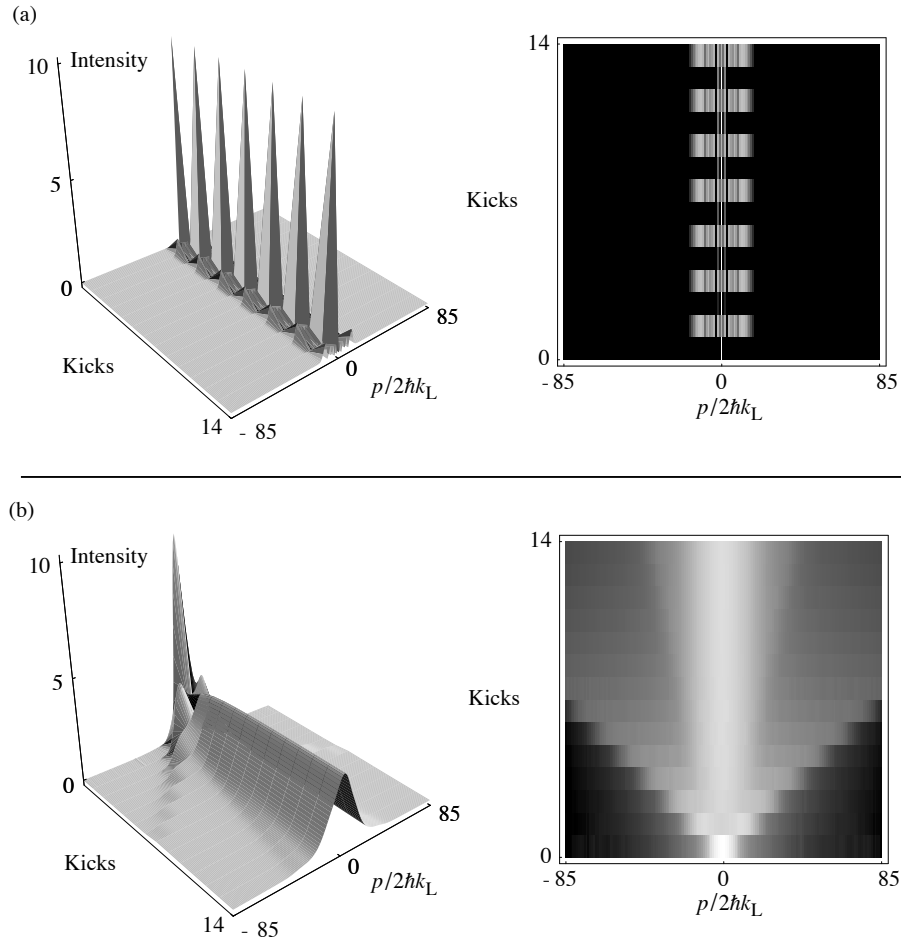


Figure 3.6: Comparison of quantum δ -kicked rotor simulations at $\hbar k = 2\pi$ and $K = 82$. Part (a) shows a simulation with a plane wave ($p_0 = 0$) initial condition. This case shows quantum “antiresonance” behavior characterized by period-2 recurrences and a lack of ballistic peaks. Part (b) shows a simulation for the same parameters, but beginning with a near-Gaussian initial momentum distribution of width $\sigma_{p_0}/2\hbar k_L = 4$. The near-Gaussian initial condition is the result a time-of-flight measurement of the initial conditions in our setup and is similar to that used in the experiments. In this case, ballistic peaks are visible and the overall dynamics are more complicated. Clearly, the antiresonance behavior is a special case that occurs for a specific momentum class. The intensity scale for the surface plots (left) is linear while the intensity scale for the density plots (right) is logarithmic.

a static, rather than accelerating, momentum distribution [Moore95; Bharucha99]. In those experiments, no difference was observed between the behavior of the system at $\hbar k = 2\pi$ and




Figure 3.7: Momentum distribution evolution at the quantum resonance: experiment. Parts (a-c) are near $\bar{k} = 2\pi$, with $K = 82$ and (d-f) near $\bar{k} = 4\pi$, with $K = 184$. While the density plots (a,d) are shown with a logarithmic intensity scale, the surface plots are shown with both logarithmic (b,e) and linear (c,f) intensity scales to highlight different features. The top “ridge” of case (c) has a period-2 series of bumps that corresponds to antiresonance behavior, which is not visible in (f). The vertical scales are in arbitrary units.

4π . Our experiment has an order of magnitude improvement in the signal-to-noise ratio over the sodium experiment, and we are able to resolve some qualitative differences between the two cases. Moreover, we can see that some part of the distribution does exhibit ballistic




Figure 3.8: Quantum simulations of momentum distribution evolution, for the same cases as the data shown in Fig. 3.7. The two cases are (a-c) near $\bar{k} = 2\pi$, with $K = 82$ and 600 ns pulses and (d-f) near $\bar{k} = 4\pi$, with $K = 184$ and 295 ns square pulses. The initial condition for both cases was a Gaussian with width $\sigma_{p_0}/2\hbar k_L = 4$. The evolution of the ballistic peaks closely matches that of the experimental data in Fig. 3.7. The remainder of the distribution undergoes motion that is primarily not resolved in the experiment. The simulations in (a-c) used 30 wavepackets, and those in (d-f) used 50.

motion.

Besides our initial distribution, there are several other differences between our sys-

tem and the ideal case. If we continue to view the pulse sequence as collapsing to a single large pulse, then the motion of atoms during the total pulse time may no longer be negligible. This leads to a momentum boundary in phase space, although the proper framework for examining this boundary is in the context of bounded motion within a nonlinear resonance [Robinson96]. A second complication is that our pulse period T (defined from the beginning of one kick to the beginning of the next one) is not exactly equal to the time *between* kicks. Recall that \tilde{k} is proportional to T and it is only the factor of \tilde{k} in the free-drift period that creates the resonance. A consequence of these effects is that there is some subtlety in comparing our experiments with theory based on δ -kicks. The most straightforward method of comparison is to perform numerical quantum simulations. The finite-pulse effects can be accounted for by directly incorporating them into simulations that take experimental parameters such as the pulse length. These simulations are used for unambiguous comparison of theory and experiment in the remainder of this section.

The experiments themselves are very similar to those that we have already presented, where the physics that we study results from a wise set of parameter choices. A pulse period of $T = 60.5 \mu\text{s}$ gives $\tilde{k} \approx 2\pi$, and $121 \mu\text{s}$ corresponds to $\tilde{k} \approx 4\pi$. Our definition of the stochasticity parameter tells us that $K \propto T$ (§1.4.3). The pulse periods used here are quite large as compared to our standard value of $20 \mu\text{s}$, and this allows us access to rather high values of K with our limited laser power. The values of K that we present here are for the “classical” stochasticity parameter, *not* the scaled quantum values K_q given by Eq. 3.4. The value of K_q goes to zero in the vicinity of a quantum resonance, but the value of K remains a useful measure of the total interaction strength. Our pulse durations for the experiments were roughly 300 ns (for $\tilde{k} = 4\pi$), but increased to 600 ns for the $\tilde{k} = 2\pi$ data set to increase K .

The atomic momentum distributions from the experiment are shown in Fig. 3.7. In this figure, we are able to resolve some component of the distribution which undergoes ballistic motion, and some (in the antiresonance case) that displays period-2 oscillation. The ballistic peaks are visible as the leading edge of the momentum distribution in these plots.

The peaks move rapidly into the higher momentum region where the signal to noise ratio is lower, and eventually exit our detection region altogether. There is also some curvature in this trajectory in momentum space which is due to the finite length of the pulses. Naturally, this curvature is more severe in the case with 600 ns pulses. The results of the matching set of numerical simulations is shown in Fig. 3.8. The primary difference between the experiment and simulation is the signal to noise ratio. Despite this, the agreement is excellent.

Numerical simulations conducted in the vicinity of quantum resonances have proven to be particularly sensitive to the chosen ensemble and grid resolution. We will describe these simulations in some detail because high resolution is required to eliminate numerical artifacts and ensure that the calculated distributions converge. The simulations with near-Gaussian initial conditions (Fig. 3.6b and Fig. 3.8) were started with a distribution that is very close to the initial condition in our experiment. An ensemble of 30 wavepackets (50 for some of the simulations shown in Fig. 3.8) was distributed uniformly along a unit cell in position. The momentum distribution for each packet was directly adapted from a recent time-of-flight momentum measurement of the initial condition in our experiment. These simulations were performed on a grid spanning the range $p/2\hbar k_L = \pm 256$. The grid resolution was $\Delta p/2\hbar k_L = 1/1024$ and the plots were smoothed by averaging over bins of $2\hbar k_L$. For Fig. 3.8, the pulses were modelled by square pulses, and we have found that the simulations are not highly sensitive to the exact pulse shape.

The ballistic peaks can clearly be visually resolved by looking at the distributions. It is harder to classify them by looking at the energy evolution. In this series of experiments, only a small fraction of the atoms are at states that undergo ballistic motion. The vast majority of the atoms are not in one of the “interesting” momentum classes and simply do not exhibit the usual quantum resonance behavior. Colder initial conditions will increase the number of atoms at $p_0 = 0$ and enable a more detailed study of these dynamics. These experiments are not the first observation of quantum resonances, however, it is the first case that we know of where the ballistic component has been resolved. Additional details about this series of experiments are presented in [Oskay00]. Let us note in closing that there is

some related ongoing work in an atom-optics experiment similar to ours. Motion near the quantum resonance with an applied acceleration has been suggested as a possible atom-optic element [Oberthaler99; Godun00].

3.5 Decoherence and the Effects of Noise

Coherent effects such as dynamical localization occur because of quantum interference. Normally, quantum effects can only be observed in highly isolated systems where coherent evolution can proceed over timescales long enough that the quantum dynamics differ measurably from the classical dynamics. The destruction of quantum interferences is known as *decoherence* and is believed to be necessary to reconcile quantum dynamics with the classical behavior of the everyday world [Zurek85]. Macroscopic objects tend to interact with their environment, and this interaction may be sufficient to destroy the fragile quantum coherences.

Usually when we discuss the classical limit of a system we wish to identify some sort of quantum-classical correspondence. That is, we expect that in some manner it should be possible to derive the behavior of the everyday world from quantum mechanics. In some simple models the classical limit might be obtained by examining the limit $\hbar \rightarrow 0$. As we have seen in our experiments so far, the *effective* value of \hbar can be changed by manipulating the scale of an experiment. Mathematical license notwithstanding, it is important to point out that as small as we can make it, $\hbar \neq 0$. The value of \hbar is already infinitesimally small when compared to the action scale of any macroscopic object. Nonetheless, simple quantum models predict that even macroscopic objects should behave in a much less classical manner than we observe [Zurek94]. It is precisely this problem that suggests that additional mechanisms are responsible for classical motion. Such a mechanism is the introduction of noise and dissipation by coupling the system to its environment, which acts to suppress quantum coherences. Decoherence in the context of a classically chaotic system is particularly interesting because of the vastly different behavior in the classical and quantum cases. In fact, it is reasonable to ask how classical chaos could be possible since every physical system that exhibits chaos is fundamentally (at the microscopic level) a quantum system. Quantum-classical correspondence in the context of chaos is a topic of general importance in physics. As such, the effects of noise and dissipation on the classically chaotic systems has been an active area of theoretical research [Shepelyansky83; Ott84; Cohen91; Shiokawa95; Cohen99; Paz01].

Although the majority of research in this field has been theoretical, there have been several experiments that address decoherence in classically chaotic systems. Key among these results have been experiments in the context of Rydberg atom ionization. We have already mentioned that these experiments have been used to observe quantum localization phenomena (§1.5). It has also been shown that the addition of noise to the periodic driving in this system brings the ionization threshold close to that predicted by classical models [Arndt91; Bayfield91; Sirko93]. Experiments conducted in condensed matter systems are usually strongly coupled to their environment through (if nothing else) finite temperature effects, which have been studied in a nonlinear system [Clarke95]. Finally, the destruction of localization by noise has been observed in the optical analogy of the kicked rotor system [Fischer00b]. There have also been some important experiments with linear systems. One system that has the advantage of simplicity is a perturbed atom interferometer [Chapman95; Kokorowski01], where spontaneous scattering introduces dissipation into the system. There have been experiments with Rydberg atoms coupled to microwave cavities [Brun96]. Trapped ions are a promising system for quantum computation and there have been studies of the decoherence of motional Schrödinger cat states [Monroe96]. Recently a “decoherence-free” quantum memory has been demonstrated that encodes a quantum bit between the states of two ions, and is robust against applied noise [Kielinski01].

One of the most important and detailed topics of study in our experiments has been the effects of noise on quantum dynamics. The fact that our system is sufficiently isolated to observe dynamical localization means that we can add noise and dissipation in a *controlled* manner. We then hope to observe noise-induced delocalization. The clear differences between the quantum and classical dynamics in this system give us a high degree of sensitivity. Besides being able to measure the energy growth of our system, we are able to directly measure our momentum wavefunctions. Both of these measures give us information about the nature of the dynamics in the system. This is an advantage over the studies with Rydberg atoms and the condensed matter experiments, in which it is not possible to observe the internal dynamics to such a degree. In the previous generation sodium-based

experiments in our laboratory, preliminary evidence was observed for amplitude and phase noise-induced delocalization [Robinson95a]. These experiments were not pursued in detail at that time because of the significant momentum-boundary effects that are present with sodium atoms. However, it was one of our first projects with the cesium-based experiment [Klappauf98b]. In the studies presented in this chapter we performed kicked rotor experiments with two types of perturbations to the system. The first perturbation is a controlled form of noise that is applied to the kick sequence. The other perturbation is dissipation applied by means of spontaneous emission. There are several important issues that we wish to address by means of these experiments. If possible, we would like to see how sensitive the system is to different types of noise. We would also like to ask if noise is sufficient to drive a system to its classical limit, and if so, how much noise is required.

We must be careful because the dynamics in the classical limit may not be obvious. When we apply noise to the kicked rotor experiment, it modifies the classical dynamics as well as the quantum dynamics, and we must take this into account. Our first series of experiments was largely qualitative, and we did not address the correspondence issue in detail. It is certainly not the case that there is a simple noise threshold above which a system behaves classically. We have a complex system where classical chaos is suppressed by quantum interferences, which are in turn destroyed by decohering processes. These competing processes each have a characteristic time scale. Dynamical localization sets in after the quantum break time, and phase coherences decay at a rate that depends upon the nature of the perturbation. For weak driving where phase coherence is lost on timescales long compared to the break time, we expect to see quantum behavior. The long-time limit of a system with weak noise may be somewhat different from a truly isolated quantum system. If quantum coherences are broken on a timescale comparable to the break time, there is little hope of seeing nonclassical behavior [Cohen91; Delande01].

We introduced amplitude noise to our experiment by replacing the kick strength K with by a step-dependent value $K + \delta K_n$. The pulse intensities are uniformly and randomly distributed in an interval centered on the zero-noise intensity level. The deviation δK_n is

then randomly chosen to be within an interval $\pm\frac{1}{2}\delta K_{p-p}$, where K_{p-p} is the peak-to-peak deviation of the kick strength. We specify the noise intensity as the ratio K_{p-p}/K . This type of noise does not substantially change the dynamics from a quantum optics point of view, in the sense that the momentum ladder is preserved. Moreover, this type of noise is (in principle) reversible.

To introduce spontaneous emission events, we leaked a small amount of near-resonant light into the chamber during the kicking sequence. We did this by turning on the molasses beams at a very low intensity while they were at their “cooling” detuning value (39 GHz red of the cycling transition). Spontaneous emission is a dissipative process, and the nature of the interaction is very different from the simple modulation of the interaction potential. Spontaneous scattering can occur in any direction, and leads to momentum diffusion (recoil heating) that is not confined to the axis of the standing wave. The momentum ladder that simplifies our analysis in other cases is broken, and the dissipative interaction is necessarily irreversible. The level of this perturbation is specified by the probability of a spontaneous scattering event per kick period (per atom). Unlike the amplitude noise which is explicitly programmed by the computer, the spontaneous emission rate figures that we cite are based upon the intensity and detuning of the light. The scattering rates were calculated under the assumption that the molasses light is uniform with all polarizations (so that the distribution of magnetic sublevel population is not important). The saturation intensity in this case is $I_{\text{sat}} = 2.70 \text{ mW/cm}^2$. We used molasses intensities up to $246 \mu\text{W/cm}^2 \pm 20\%$, and the total scattering rates are therefore only determined to this precision.

It is possible to increase the spontaneous emission rate from the standing wave by simply detuning closer to resonance. This method was employed by Christensen’s group [Ammann98] in an experiment to look for decoherence effects. Many of the noise sources and deviations from an ideal two-level quantum system become more severe as the detuning is brought closer to resonance. The effects of spontaneous emission can seriously complicate the analysis of the system, especially considering that it can change the internal state of the

atom. Beyond this, classical boundary effects and the effects (at their parameters) of the stochastic dipole force [Gordon80] make their results difficult to interpret [Habib98]. Our experiment with molasses light clearly separates all noise sources inherent in the standing wave (especially those that depend upon the detuning) from the dissipation.

For these experiments we specialized to a pulse period of $T = 20 \mu\text{s}$, so that $\hbar k = 2.08$. The detuning was again 6.12 GHz red of resonance, and we used the same 283 ns (FWHM) pulses that we have described earlier. We picked values of K that gave us the best exponentially-localized momentum distributions in the zero-noise case. These cases are near the first true minimum in the $D(K)$ curve, at $K_q \approx 10$ (see Fig. 3.5). For the amplitude noise data presented in this section, $K = 12.8$, and $K = 11.9$ for the spontaneous emission case. At our value of $\hbar k$, the quantum scaling of K (Eq. 3.4) leads to a shift of about 20% in the locations of the peaks of the diffusion curve.

The effects of amplitude noise and spontaneous emission on the energy growth of the kicked rotor are shown in Fig. 3.9. The zero-added noise cases exhibit the basic trends that we expect for dynamical localization. The energy growth is only initially in agreement with classical expectations, and after a short time the energy growth drops dramatically. As we add amplitude noise or dissipation, we see these curves bend towards the classical expectations. The amplitude noise data were averaged over four distinct realizations of the noise (the error bars are purely statistical). In the case of spontaneous emission, it is important to identify the role of recoil heating, which could potentially mimic the effects of decoherence by increasing the energy of the atomic sample. To this end, we measured the effect of turning on the optical molasses beams *after* the interaction sequence instead of during it. The molasses light was introduced at the same intensity and for the same duration as that which we used during the kick sequence. The results of this measurement are shown as an inset to Fig. 3.9. Because we have used optical molasses, as the intensity of the light is increased, it actually has the effect of slightly *cooling* the distribution. This is clear evidence that our source of energy growth is the destruction of localization, rather than recoil heating. As we have discussed earlier, the absolute values of the energy values




Figure 3.9: Noise and dissipation effects on energy evolution. The energy $E = \langle (p/2\hbar k_L)^2 \rangle / 2$ is plotted as a function of time for several levels of (a) amplitude noise and (b) spontaneous emission. The amplitude noise levels are 0%, 25%, 50%, and 62.5%, monotonically increasing in energy. The spontaneous emission rates are 0%, 1.2%, 5.0%, and 13% (circles, filled triangles, open triangles, and diamonds) probability per kick period, representing intensities up to $246 \mu\text{W}/\text{cm}^2$. The solid lines interpolating the data points are fits to a simple model ([Klappauf98b]). The other two lines are classical simulations of the δ -kicked rotor (solid) and the square-pulse kicked rotor (dashed). The inset (b) shows the effects of exposing the system to the near-resonant light *after* the sequence of kicks.

may have significant systematic errors. Furthermore, the data sets shown in Fig. 3.9 are the last case that we show that still had the residual $F_g = 3$ ground state population. Again, this leads to an overall reduction in the energy values, but does not affect the physics that we are studying.

The momentum distribution evolution for various levels of spontaneous emission is shown in Fig. 3.10. Without added noise, we see dynamical localization, and with increasing noise level, the distributions increasingly depart from this. For the higher cases of sponta-




Figure 3.10: Spontaneous emission effects on momentum distribution evolution. The four levels of spontaneous emission shown, 0%, 1.2%, 5.3%, and 13%, are near those that the energy growth is shown for in Fig. 3.9. The distributions for higher levels of dissipation appear to be roughly Gaussian in character, which suggests that the dynamics are tending towards classical.

neous emission probability, the distributions appear rather rounded and are suggestive of the Gaussian distributions that we expect in classical dynamics.

Qualitatively, the effects of dissipation and amplitude noise on our system are not very different. In both cases we see the destruction of localization in the energy evolution of the system. The momentum distributions in the case of amplitude noise are qualitatively similar to those that we see for applied dissipation. We have studied the effects of amplitude noise on the kicked rotor system in depth, and will present these results (including the momentum distribution evolution) in the next section (§3.6). We have also observed the destruction of quantum effects by other types of noise that we have not studied in detail. Phase noise, spontaneous emission from an independent far-detuned beam of light, and timing noise in the period of the kicking sequence have each been observed to decohere

our systems. As an example of one of these cases, the effects of timing noise are shown in §3.7. We have not observed substantial qualitative differences between the effects of these noise sources, however we have not conducted the detailed study that would be necessary to compare them.

So long as we wish to study the effects of small amounts of noise, it is necessary to identify the residual noise sources in our experiments. In describing the experiment, we have detailed the phase and amplitude noise that are always present in our one-dimensional standing wave. However, the most important source of noise that we have is residual spontaneous emission that occurs from the intense light in the standing wave. The light is detuned 6.12 GHz to the red of the cycling transition, and is on for a very small fraction of the total interaction time. We have estimated the spontaneous scattering probability from the standing wave to be less than or about 1% per kick. Naturally, this depends entirely upon the intensity of the light that we are using, which is highest when we are tracing out $D(K)$ curves. The intensity for the noise cases that we have presented (thus far) were at about 1/3 of the maximum intensity. Note that this can still result in a 20% probability of spontaneous emission by the end of the experiment. Nevertheless, the system is still very sensitive to small amounts of added spontaneous emission, as we see in Fig. 3.9. Spontaneous emission has several possible effects. Beyond the decohering effects and heating, a spontaneous scattering event can result in an atom decaying to the “wrong” ground state, as we discussed earlier (§1.4.2).

Collisions could also act as a potential source of decoherence in our experiments, but are more unlikely. The highest density that we have operated at was for this series of experiments, where the total atom number density was of order 10^{11} cm^{-3} . Our initial condition has a mean atomic velocity of order 10 recoil velocities. Based on the collision cross section for cold cesium atoms [Arndt97], we estimate the maximum collision probability to be 2%/ms. For the experiments that we will describe later with velocity selection, the density is much lower (10^8 cm^{-3}) and collisions are negligible. There are also small (negligible) perturbations to the system from the stochastic dipole force, nonuniformity of the

interaction strength due to the magnetic sublevel distribution (§1.4.1), and nonlinearities due to higher order terms in the interaction potential. These effects are treated in detail by Steck [Steck01b]. Each of these residual noise sources acts as a minor offset to the location of the “zero” noise case. Furthermore, these effects collectively contribute to the residual energy growth that our system exhibits after localization sets in. Generally speaking, we cannot quantitatively correct for these sorts of effects. There are several systematic effects that *can* be accounted for, however, as we will see shortly.

3.6 The Classical Limit

Thus far we have described some initial experiments in which we observed the effects of noise and dissipation on the quantum kicked rotor system. The quantum dynamics are sensitive to these sorts of perturbations and we have seen that they can destroy dynamical localization. While the atomic momentum distributions that we have observed are *suggestive* of classical dynamics, we have not established quantitative correspondence between our experiment and a classical description of our experiment. To this end, we have performed a detailed study of amplitude-noise effects on dynamical localization [Milner00; Steck00]. There are two essential parts of this study: a more detailed experimental study of quantum dynamics in the presence of amplitude noise, and a detailed classical model that we compare with the experiment. The experiment is performed with a similar set of parameters to the amplitude noise experiments that we have already described. The essential difference between this and the previous data set is the larger range of both noise level and stochasticity parameter that we present here. In this section we describe many of the complications in accurately comparing our experiment to a computational model.

3.6.1 Amplitude noise effects on kicked rotors

Because we wish to compare our experimental results with a classical model, it is important to understand the effects that amplitude noise has on a *classical* kicked rotor. The complex structures that occur in the standard map and the effects of correlations that we




Figure 3.11: The classical diffusion rate is plotted in the presence of increasing levels of amplitude noise. These values of $D(K)$ are given by Eq. 3.6, with the uniform amplitude noise described in the text. The curve with 0% applied amplitude noise (heavy solid) is a limiting value that matches the curves that we have encountered earlier. As the peak-to-peak deviation of the noise is increased to 50% (dashed), 100% (dotted), and 200% noise (dot-dash), the short-term correlations are gradually washed out by the noise.

have seen suggest that this system will be affected by the introduction of amplitude noise, and so ultimately we must compare our experiment to a noise-driven rotor. The realization of the amplitude noise is the same uniform distribution that we used in the preliminary experiments. The probability distribution of kick strengths is given by $P(\delta K) = 1/\delta K_{p-p}$ if $|\delta K| \leq \frac{1}{2}\delta K_{p-p}$, and zero otherwise. The noise modifies the classical correlations since the value of K for each kick is now independent. A generalization of Eq. 1.15 (from [Rechester81]) that allows for randomly distributed kick strengths $K_n = K + \delta K_n$ can be considered by averaging each kick strength with its probability. The generalized diffusion curve in the presence of noise is

$$D(K) = \frac{K^2 + \text{Var}(\delta K)}{4} + \frac{K^2}{2} (-\mathcal{J}_2(K) - \mathcal{J}_1^2(K) + \mathcal{J}_2^2(K) + \mathcal{J}_3^2(K)), \quad (3.6)$$

where $\text{Var}(\delta K)$ is the variance of $P(\delta K)$, which is $\text{Var}(\delta K) = (\delta K_{p-p})^2/12$ for uniformly distributed noise [Steck00]. The Bessel functions in the noiseless case have been replaced

with

$$\mathcal{J}_n(K) \equiv \int_{-\infty}^{\infty} P(\delta K) J_n(K + \delta K) d(\delta K), \quad (3.7)$$

which is much like convolving the Bessel functions with the noise probability distribution. Diffusion curves given by Eq. 3.6 for several values of amplitude noise are shown in Fig. 3.11. For low levels of amplitude noise, the oscillations in the diffusion curve are damped, and the curve inches towards the quasilinear limit (which we derived from the assumption of no correlation). For higher levels of noise, the quasilinear diffusion rate (identified here as $D_{\text{ql}}(K) = [K^2 + \text{Var}(\delta K)]/4$) itself begins to climb noticeably. The maximum level of amplitude noise that can be applied in our realization of amplitude noise (without going negative) is 200%, where $\delta K_{\text{p-p}} = 2K$, twice the mean amplitude. We note that at this maximum value, $\text{Var}(\delta K)/4 = K^2/3$. This indicates that although the noise is sufficiently strong to destroy correlations, it does not dominate the energy growth.

The effects of amplitude noise on the quantum correlations can be treated in a similar manner by generalizing Eq. 3.5. The resulting quantum diffusion curve with amplitude noise is described by

$$D(K, \hbar) = \frac{K^2 + \text{Var}(\delta K)}{4} + \frac{K^2}{2} (-\mathcal{Q}_2(K_q) - \mathcal{Q}_1^2(K_q) + \mathcal{Q}_2^2(K_q) + \mathcal{Q}_3^2(K_q)). \quad (3.8)$$

The Bessel functions here have been replaced by

$$\mathcal{Q}_n(K_q) \equiv \int_{-\infty}^{\infty} P(\delta K) J_n(K_q + \delta K_q) d(\delta K), \quad (3.9)$$

where δK_q is the scaled quantum kick strength deviation, $\delta K_q = \delta K \text{sinc}(\hbar/2)$. The form of these expressions is almost exactly that of the classical expressions, and we can get a good deal of intuition about the effects of noise on quantum correlations by examining the classical case. We then expect the quantum short-term correlations to wash out with increasing levels of amplitude noise. The primary difference between the classical and quantum diffusion curves is the location of the peaks, which is shifted by about 20% for $\hbar = 2.08$. It is then reasonable to expect that we will not see quantum-classical correspondence for noise levels that are low enough that we can resolve the locations of the peaks.

3.6.2 Systematic effects and the classical model

With some conception of how the dynamics will change as we add noise, we can now begin to consider the classical model that will be compared with our experiment. Our experiment has several deviations from an ideal delta-kicked rotor. The most important deviations are the nonzero length of the pulses and the shortcomings in the detection system. There are also some other subtle issues in the total timescale of the experiment. We have already discussed finite-pulse effects in some detail, and we have mentioned some of the systematic issues in our detection stage. Beyond these issues, the width of our initial condition influences the final distributions that we observe. One can imagine directly correcting the measured distributions for some of the systematic effects in the imaging systems. However, it may not be possible to correct for more complex problems such as the pulse width. It is much simpler to directly incorporate each of these effects into a classical simulation that we can compare with our experimental data. The computational cost of incorporating each of these effects into an accurate and useful quantum simulation is high, and possibly prohibitive.

The classical Monte Carlo simulations that we will discuss here were written by Daniel Steck, and are described elsewhere in greater detail than will be presented here. [Steck00]. For these calculations, the trajectories of an ensemble of $2 \cdot 10^5$ particles were calculated over the course of the interaction. Each particle experiences an independent sequence of randomly chosen kick amplitudes so that we effectively average over many possible realizations of the noise. The initial distribution is taken directly from an experimentally measured momentum distribution, much like in the simulations performed for the quantum resonance experiment.

We have already seen how the effects of finite-length pulses reduces the effective stochasticity parameter as a function of momentum (§3.2). The square pulses that we described earlier are sufficient to predict many properties of the actual pulses that we use. However, it is important to model the pulses as accurately as possible in order to correctly account for motion that occurs at high momenta. We recorded the profile of a typical “kicking” pulse with a fast photodiode and our digitizing oscilloscope, in the usual configuration




Figure 3.12: The profile of a single optical pulse used in our kicked rotor experiments. The experimental signal (dashed) was recorded by a photodiode and oscilloscope, and compares very well to a model function (solid) that we use to describe the pulse in the classical simulations (Eq. 3.10).

that we use to monitor the interaction beam. The profile that we have measured fits well to an empirical model function

$$f(t) = \frac{1}{2\eta_{\text{erf}}} \left[\text{erf} \left(\frac{(t - t_1)\sqrt{\pi}}{\delta t_1} \right) - \text{erf} \left(\frac{(t - t_2)\sqrt{\pi}}{\delta t_2} \right) \right], \quad (3.10)$$

which is defined in terms of the error function $\text{erf}(x) \equiv 2/\sqrt{\pi} \int_0^x \exp(-s^2) ds$. The full width at half maximum of the pulse is $t_2 - t_1 = 295$ ns. The rise and fall times of the pulse are $\delta t_1 = 67$ ns and $\delta t_2 = 72$ ns, respectively, and η_{erf} is a normalization constant. This pulse function and the experimental pulse that it models are shown in Fig. 3.12. The effects of this pulse are incorporated into the simulations by directly integrating the equations of motion over the pulse profile.

Now we come to some of the systematic effects associated with our detection methods. The most important of these effects is the nonuniformity of the molasses beams that we use to image the atoms. We freeze the atoms in the molasses during the imaging time, and we have assumed thus far that the spontaneous scattering rate of each atom is independent of its position. Unfortunately, the molasses region is formed by the intersection of several

Gaussian beams with an intensity profile that decreases towards the edges of the beams. The result is that atomic populations in the higher momentum regions appear to be reduced because the atoms there do not fluoresce as brightly. The extent of the falloff was measured by examining a momentum distribution that was known to be exponentially localized at several drift times. This procedure is more precise than estimating the fluorescence rates from the laser parameters. A similar drop-off in population can be applied at the detection phase of the classical simulations. A second detection issue is the finite size of our detection region. Typically we use a 15 ms free drift period that allows us to detect atoms at momenta within $\pm 80 \cdot 2\hbar k_L$. Although the falloff of the molasses beam intensities becomes worse outside this region, we are more firmly limited by the region that we image with our CCD camera. For long time evolution, especially in cases with high levels of amplitude noise, our detection window may truncate the momentum distribution. The most severe effect of the window is a systematic reduction in the energy for large energy values— it is only at high energy values that a significant population leaves our detection region.

There are several other effects that we must account for, mostly concerned with the finite extent of the experiment in both space and time. The standing wave that the atoms interact with has a beam waist of about 1.5 mm, which is an order of magnitude larger than the characteristic spatial extent of the atoms collected in the MOT. However, the kick strength that each atom experiences depends (weakly) upon its transverse spatial location in this beam profile. The transverse distribution of the initial position and momentum values is modeled by the (transverse) size and position of the atoms released from the MOT, as determined by temperature measurements. During the course of the experiment (which may last several milliseconds), the atoms fall gravitationally through the transverse profile of the beam. All of these effects can be taken directly into account in the simulations by assigning a time and space-dependent stochasticity parameter for each atom. The chief effect of these complications is that the effective value of K is smeared by around 5%. This is not a tremendous problem for us; we are clearly able to resolve features in the diffusion curve (as in Fig. 3.4). Incidentally, this is the most important reason for maintaining a relatively large

interaction beam spot size. Christensen’s group was unable to resolve any deviation from the quasilinear limit in the diffusion curve [Ammann98], presumably because of this very effect. We have also introduced the ambiguity in the definition of the free drift time that results from atomic motion *during* the interaction time (§2.7). It is fairly straightforward to eliminate the ambiguity by simulating the free expansion that occurs during and after the interaction. We might also be concerned about the initial spatial size of the distribution and how it affects the final distribution that we measure. As we have discussed, it is not generally practical to separate this component of the distribution except in the case of temperature measurements (§2.8). The initial momentum distribution that is used in the simulation is derived from a free-expansion measurement and already “includes” the initial spatial extent of the distribution (convolved with the true momentum distribution). We then note that the simulation *implicitly* accounts for the convolution with the initial spatial distribution.

A final complication of our experiments results from the behavior of our Princeton Instruments CCD camera. As we have said, the background level and absolute gain of this camera tend to fluctuate. We employ background subtraction to take out the largest and slowest drifts. What remains after we subtract the background image from each camera image is primarily the atomic distribution that we are after, but there is usually some residual offset in the distribution. More troubling than this is the fact that the total intensity at some points in the distribution is negative. This can occur when a pixel in the background image is more intense than a pixel in our picture of the atoms. We integrate the pictures along one axis to create the one dimensional momentum distribution. Usually, tens of the points in the 1-D array (out of 510) remain negative. We could simply add a constant to all of the values to keep them all greater than zero. This would pose a problem for normalization and subsequent analysis of the distribution, which can be very sensitive to this variation. The energy values are particularly sensitive to the absolute offset of the distribution at high momentum. Rather than adding a constant, we employ a procedure that is designed to maximize the reproducibility of our data. We average together the values of the lowest 40 points in our distribution and consider their average value to be the “zero”




Figure 3.13: Illustration of systematic corrections to classical simulations. The system energy is plotted as a function of time for classical kicked rotor simulations. The system parameters are typical of our experiments, $K = 11.2$ with 100% amplitude noise. First, the noise-driven δ -kicked rotor is shown without any of the systematic effects (solid), and appears as a straight line. As we cumulatively add in the systematic effects for typical operating parameters, the energies at later times gradually decrease. The effects shown here are the finite pulse duration (dashed), the molasses beam intensity profile (long dot-dash), clipping at the edges of the detection region (dotted), motion in the transverse profile of the interaction beam (long dashes), drift time ambiguity (dot-dash), and background-subtraction bias (thin solid).

level of our distribution, and discard values below that threshold. For distributions with low energies, the values that we average together are typically at the edge of the distribution where there essentially no atoms. This provides an accurate estimate of the background level of the camera during that particular shot. When the energy is higher and the distribution reaches the edge of the imaging window, this procedure still helps to increase reproducibility. However, in this case it also has the effect of negatively biasing the entire distribution. Much like the other cases that we have discussed, this effect can be directly accounted for by using the similar analysis methods in the case of the simulation.

We have now described a number of systematic corrections that must be considered if we wish to understand the classical limit of our system. The factors are most significant for higher momentum values, which we encounter at later times in the experiment and/or

when there are high levels of amplitude noise. In a classical δ -kicked rotor we expect the energy of the system to grow linearly with time. However, the correction factors reduce both the kicking strength and the detection efficiency at high momentum. The energy growth in the simulation is curtailed as illustrated in Fig. 3.13. The final curve that we see in this figure tells us the sort of momentum distribution that we expect to observe *if* the dynamics are described by a classical model.

Generally in these simulations, there are no adjustable parameters. Each of the parameters that is used in the classical model is derived from measurements in our experiment. The largest source of uncertainty in our experiments is the absolute (not relative) value of K , which is not determined better than $\pm 10\%$ for this series of experiments. In some of the simulations, the value of K was shifted slightly (within the uncertainty) for better agreement with the experiment. The particular value of K used in each simulation is indicated in the caption of each figure in which we show the results of one of these simulations.

3.6.3 Comparison of the experiment to the classical model

Let us now present the results of this experiment and compare them with the classical model. We begin by examining the behavior of the diffusion rate curve $D(K)$ in the presence of amplitude noise. We have already considered the effects of amplitude noise on both the quantum and classical kicked rotor (§3.6.1). We expect the short-term correlations in both cases to be washed out by amplitude noise. The most significant difference between the two cases is the location of the peaks, which is shifted by about 20% with our value of \hbar . We have also noted that the difference in the peak location suggests that correspondence is not possible until the short term correlations are washed out. The diffusion curves are traced out by data from our experiment and classical simulations in Fig. 3.14. As before, we do not measure $D(K)$ directly, but measure $E(K)$ at a fixed time. For this particular data set we measured the energy of each distribution after 35 kicks. Dynamical localization occurs at low noise levels and reduces the typical energies that we observe. For low levels of noise, we see that the quantum and classical dynamics are clearly separate. As we turn up




Figure 3.14: The energy as a function of K with amplitude noise: comparison of experiment (connected points) to classical simulations (dashed line). The energy is measured after 35 kicks. For 0% amplitude noise, the classical and quantum curves both display large oscillations, separated substantially by their location in K . As the noise level is increased, the correlations that create the oscillations about the quasilinear diffusion rate are damped away. At the strongest noise level shown here, the experimental curve is indistinguishable from the classical simulation. Each point represents an average over 10 experiments with independent realizations of noise. The values of K shown here are experimentally determined. The statistical errors in both K and energy for each point are similar to or smaller than the size of the dots. The arrows label two values of K at which we study the dynamics in greater detail.

the noise, we see that the energy curve approaches the classical expectation. This in itself is *not* confirmation that the system is behaving classically. We must also trace the energy evolution as a function of time. Beyond this, we must also look for correspondence between the classical and experimental momentum distributions. We will examine the behavior of the system in greater detail at two values of K , which are indicated by arrows in Fig. 3.14.

We now examine the effects of noise in a regime where dynamical localization would otherwise occur. This study was conducted near the first minimum in the diffusion curve, shown by the rightmost arrow in Fig. 3.14. We operated at $K = 11.2$, which is close to the value that we used in the earlier amplitude noise data (§3.5). The evolution of the system energy as a function of time is shown in Fig. 3.15 for various noise levels. For low levels of noise we see the effects of dynamical localization: the energy growth of the system is sharply lower than we would expect from classical predictions. Besides this, there are differences between the classical and quantum short-term correlation effects. As the noise level is increased, the energy of the corresponding classical dynamics also increases, but this change is substantially smaller than the difference between the classical and quantum curves at low noise levels. For higher noise levels (60% and above), the quantum evolution is indistinguishable from the classical expectation.

We can attempt to identify the noise level at which the quantum and classical energies diverge by plotting the energy as a function of the noise level. We compare the experimentally determined energies to the classical simulation in Fig. 3.16. The energies are measured after 50 kicks, and are averaged over 18 distinct realizations of noise in this particular data set. We can now see more directly that the energy evolution matches the classical evolution for amplitude noise levels above about 60%.

Finally, we directly examine the momentum distributions as we apply noise to the system, shown in Fig. 3.17 for $K = 11.2$. This is the most direct and complete way to compare the evolution in our experiment to classical expectations. For low levels of noise, we find that the momentum distributions in our experiment do not agree at any stage (except 0 kicks) with the classical model. For 0% applied noise we observe clear localization. By




Figure 3.15: The energy as a function of time with amplitude noise: comparison of experiment (connected points) to classical simulations (dashed line). These experiments were performed with $K = 11.2 \pm 10\%$. The corresponding classical simulations were performed with $K = 11.2$. Each experimental point represents an average over 15 experiments, each with an independent realization of the noise. The data for noise levels greater than 0% are sequentially offset by 200 in energy for the purpose of clarity.




Figure 3.16: The energy as a function of amplitude noise level: comparison of experiment (connected points) to classical simulations (dashed line). These experiments were performed with $K = 11.2 \pm 10\%$. The classical simulations were performed with $K = 10.9$. Each experimental point represents an average over 18 experiments, each with an independent realization of the noise. The energy in each case is measured after 50 kicks. The error bars which represent only the statistical variation among the 18 experiments become more visible as the amplitude noise level increases. The experiment is well described by the classical model for higher noise levels.

the time that we have applied 40% noise, the distributions match classical predictions very well, but only for short times. For higher noise levels, we observe that the correspondence between the experiment and the classical model is excellent, for the entire evolution that we are able to observe.

Let us now turn to a similar detailed study of noise in a regime of classical anomalous diffusion. This study was conducted at the first maximum in the diffusion curve, shown by the leftmost arrow in Fig. 3.14. The time evolution of the ensemble energies is shown in Fig. 3.18. The time evolution of the energy is somewhat less simple than it is in the localized case at the minimum of the diffusion curve. A possible interpretation in the low noise cases is that the quantum system diffuses rapidly before localizing. For larger levels of noise, the agreement between the experiment and the classical noise level is excellent. Again, these higher noise levels seem to begin at around 60% amplitude noise.




Figure 3.17: The momentum distribution evolution is shown for our experiment (solid) and the classical model (dashed) for several levels of amplitude noise. These experiments were performed with $K = 11.2 \pm 10\%$, and the classical simulations were performed with $K = 11.2$. This is a regime where we expect to see dynamical localization. The data shown here are the distributions from which the energy values shown in Fig. 3.15 were calculated. The distributions are shown after 0, 10, 20, 40, and 80 kicks, in order of increasing width. For 0% applied noise, the experimental distribution is clearly localized, but the classical evolution limits to a Gaussian distribution. Again, for noise levels at or above 60%, the experimental distributions closely match the classical simulation. We can see directly in this figure that the width of the classical simulation does indeed grow as amplitude noise is applied.

The energy as a function of noise level is shown in Fig. 3.19. The difference between the zero-noise behavior (after 50 kicks) in the experiment and the classical model is more significant here than in the localized case. The data from the experiment show a clear threshold for matching the classical behavior. It is worth noting that the total change in the classical energy as a function of noise is rather small and nonintuitive.

Finally, we present the momentum distribution evolution in the presence of amplitude noise, with $K = 8.4$ in Fig. 3.20. As opposed to the measurements with higher values of K , the classical distributions do not widen significantly as the noise level is increased. Once again, we see that the distributions do not match up for low levels of noise, but our experiment is described very well by a classical model for higher levels of noise.

We have seen that our experiment can be well described by a classical model for




Figure 3.18: The energy as a function of time with amplitude noise: comparison of experiment (connected points) to classical simulations (dashed line). These experiments were performed with $K = 8.4 \pm 10\%$. The corresponding classical simulations were performed with $K = 8.7$. Each experimental point represents an average over 15 experiments, each with an independent realization of the noise. The data for noise levels greater than 0% are sequentially offset by 200 in energy, as in Fig. 3.15.




Figure 3.19: The energy as a function of amplitude noise level: comparison of experiment (connected points) to classical simulations (dashed line). These experiments were performed with $K = 8.4 \pm 10\%$. The classical simulations were performed with $K = 8.4$. Each experimental point represents an average over 18 experiments, each with an independent realization of the noise. The energy in each case is measured after 50 kicks. The error bars which represent only the statistical variation among the 18 experiments become more visible as the amplitude noise level increases. The ensemble energy is well described by the classical model for higher noise levels, here above about 50%.

levels of amplitude noise above 60%. At high values of noise, this model accounts for the sensitivity to noise, the growth of the ensemble energy as a function of time, and even the detailed momentum distribution evolution. While it is not possible to confirm that the dynamics are *completely* classical, it is fair for us to say that we do not observe any deviation from classical physics over the timescales of our experiments.

3.7 Timing Noise

Let us now briefly consider one additional variety of noise. In the kicked rotor experiments that we have described thus far, the period T between kicks is a constant. Since the experimental kicking sequence is controlled by the computer, it is relatively easy to change the period between subsequent kicks. Suppose that we replace the constant T with a step-dependent period $T_n = T + \delta T_n$, where δT_n is a step-dependent perturbation to the period. As in the case of the amplitude noise, we consider a uniform probability distribution




Figure 3.20: The momentum distribution evolution is shown for our experiment (solid) and the classical model (dashed) for several levels of amplitude noise. These experiments were performed with $K = 8.4 \pm 10\%$, and the classical simulations were performed with $K = 8.7$. This is a regime where the corresponding classical dynamics exhibit anomalous diffusion. The data shown here are the distributions from which the energy values shown in Fig. 3.15 were calculated. The distributions are shown after 0, 10, 20, 40, and 80 kicks, in order of increasing width. For 0% applied noise, the experimental distribution is similar to the broad-shouldered distribution that we have observed earlier near the peaks of the diffusion curve (Fig. 3.5). Again, for noise levels at or above 60%, the experimental distributions closely match the classical simulation.

$P(\delta T_n) = 1/\delta T_{p-p}$, if $|\delta T_n| \leq \frac{1}{2}\delta T_{p-p}$, and zero otherwise. The peak-to-peak range δT_{p-p} is specified as a percentage of T , so that 100% noise means that a kick may be displaced by any value in the range $\pm T/2$. In the experiment, we would like to avoid approaching 100% noise because the pulses (of finite duration) may overlap. The system, however, is very sensitive to timing noise, and we did not find it interesting to study noise levels higher than 20%.

Our realization of timing noise is a *cumulative* one. The time that each kick occurs is $T + \delta T_n$ after the *previous* kick, rather than simply being δT_n displaced from the time that the kick would occur if there were no timing noise. This distinction is very important because kicks later in the sequence can potentially become very far displaced from their original times. The amplitude noise that we studied was a noise applied only to the value of K . In the case of timing noise, we must note that the variable T enters into both the




Figure 3.21: Timing noise effects on energy evolution. The energy $E = \langle (p/2\hbar k_L)^2 \rangle / 2$ is plotted as a function of time for several levels of timing noise, near $K = 8$. The noise levels are 0%, 0.5%, 1%, 2%, 3%, 10%, and 20%, monotonically increasing in energy. The data points are collected every five kicks, and the lines connect these points to guide the eye. Each point represents an average over 20 measurements, each with an independent realization of the timing noise. The error bars on each point are purely statistical.

definitions of K and \tilde{k} , and the (constant) average value $\langle T \rangle$ must be used for both of these variables. Given this, the standard map changes under the influence of timing noise to be

$$\begin{aligned} p_{n+1} &= p_n + K \sin x_n \\ x_{n+1} &= x_n + (1 + \epsilon_n)p_{n+1}, \end{aligned} \quad (3.11)$$

where $\epsilon_n \equiv \delta T_n / T$. The noise enters the equations of motion in a different manner than does amplitude noise. One immediate effect of this is that even a small value of ϵ_n can cause large changes to the system if p is large. Intuitively, the addition of timing noise should have a substantial effect on the corresponding classical dynamics, since the structures in phase space that we have studied are only visible through periodic (stroboscopic) sampling. Loosely




Figure 3.22: Timing noise effects on on $E(K)$ curves. The energy $E = \langle (p/2\hbar k_L)^2 \rangle / 2$ is plotted as a function of the stochasticity parameter K for several levels of timing noise. The noise levels are 0%, 0.5%, 1%, 2%, 3%, 5%, 8%, 10%, and 20%, monotonically increasing in energy (near $K = 8$). This data set was collected after 35 kicks, and each point represents an average over 20 measurements, each with an independent realization of the timing noise. The error bars on each point are purely statistical.

speaking, fairly low levels of timing noise are needed to destroy dynamical localization. The noise that we apply is only weak in the sense that we do not approach the maximum value of noise that could potentially be applied to the system. This is not a rigorous statement about the sensitivity of the system to different realizations of noise. A direct comparison would be complicated by the fact that this noise realization is cumulative, which is not the




Figure 3.23: Energy growth as a function of timing noise level. The energy $E = \langle (p/2\hbar k_L)^2 \rangle / 2$ is plotted as a function of the timing noise level for three levels of the stochasticity parameter K . The values of K are 6.8 (blue), 7.85 (red), and 8.65 (green). This data set was collected after 35 kicks, and each point represents an average over 20 measurements, each with an independent realization of the timing noise. The error bars on each point are purely statistical.

case for the amplitude noise experiments.

The experiments in timing noise were conducted after we finished our detailed study of amplitude noise. The effects of timing noise on the kicked rotor system are qualitatively similar to those of amplitude noise, and we chose to move on to other experiments rather than study the effects of timing noise in more detail. Nevertheless, we did collect some interesting data, which we will present here. At the time of these experiments we were still developing our 3D lattice cooling techniques (§4.2), and we were not yet able to fully null out the magnetic fields in the chamber. As a consequence of this, we were not able to achieve long storage times in the lattice and our ultimate temperatures were near 700 nK.

(We will describe 3D lattice cooling in detail later.) The signal-to-noise ratio for these data sets is comparatively high, as is apparent in Fig. 3.21. This is due in part to the relatively cold initial conditions that we achieved through the lattice cooling. There is also a sense in which the very nature of the timing noise sequence produces less jitter in the output signals than does the amplitude noise. One way to see this is to compare the size of the statistical error bars on the amplitude and timing noise data sets.

In Fig. 3.21 we observe the effects of timing noise on the system energy growth as a function of time. For as little as 0.5% added timing noise, there is a significant change in the long-time diffusion rate. The energy growth begins to saturate when we add timing noise at about the 3% level. This is in contrast to the amplitude noise case, where we had to make severe changes to the pulse sequence before the energy growth saturated. This saturation phenomenon is to some degree an effect of the particular value of K that we have chosen, in this case near $K = 8$. In Fig. 3.22 we present an $E(K)$ curve where we have added timing noise. For our value of K (in the first local minimum of the $E(K)$ curve), most of the energy growth occurs by the 3% noise level. Near the first local maximum of the curve (near $K = 2\pi$), the energy growth rises and then falls as a function of the noise level, indicating that the noise becomes strong enough to break the correlations that give rise to that maximum. Indeed, one of the major effects of the noise is to wash out the oscillations in the curve, just as we saw in the amplitude noise case. One additional feature of this figure is the energy growth as a function of noise level near $K = 3$. At this value of K , the diffusion rate should be given approximately by the quasilinear value because correlations do not play a major role. Finally, in Fig. 3.23 we plot (for several values of K) curves that directly show the energy growth as a function of the noise level. The energy of the system is universally increased by the addition of noise. But, beyond a certain saturation point, higher levels of timing noise can either increase or decrease the energy of the system, depending if we began at the peak or the valley of the diffusion curve. At a certain point in the middle, the system is not sensitive to additional noise. These results are all qualitatively similar to those that we saw in the case of amplitude noise, but the improved signal to noise ratio and reduced

scatter of these data sets compliments the earlier data well.

Chapter 4

Quantum state preparation

4.1 Overview

In most of the experiments that we have described thus far, the system under study has been one that exhibits globally chaotic motion. The initial conditions that we have used for these experiments have been fairly cold, with temperatures on the order of $10 \mu\text{K}$. The atomic distribution is also reasonably compact in its spatial extent with a characteristic width of $\sigma_x = 0.15 \text{ mm}$. These values seem rather small from a macroscopic point of view but are actually fairly large when compared with the typical sizes of structures in phase space. Considering that a unit cell of the kicked rotor phase space spans 2π of x and p in scaled units (§1.4.3), we find that our distribution spans (roughly) a thousand unit cells in position and one in momentum. For some studies of quantum dynamics, the distributions that we can create with the MOT are simply too broad. We have already seen in the example of the quantum resonance that the behavior of a quantum system can depend qualitatively on the initial condition (§3.4). Besides effects of this nature, we are also interested in studying quantum transport in *mixed* phase space, which consists of significant regions of both regular and chaotic dynamics. One such system that we will introduce is the potential given by an amplitude modulated standing wave. There are small stable islands in this system into which we would like to “load” our initial condition. For such specialized tasks as these, we have developed a method of quantum state preparation. The first part of the procedure involves a three-dimensional, optical lattice that we use in a final stage of cooling after we turn off the MOT. After this, we use velocity-sensitive stimulated Raman transitions and subsequent manipulation in a one-dimensional standing wave to center an atomic sample

directly on the feature of phase space that we wish to study. All together, these procedures represent the substantial changes to the experiment that we have made since the kicked rotor experiments that have already been described. We will describe these methods in detail in the remainder of this chapter.

4.2 Three Dimensional Lattice Cooling

The temperatures that we can easily achieve in our MOT, including the simple polarization gradient cooling (PGC) period, are not much better than $10 \mu\text{K}$. This temperature corresponds to a momentum width of $\sigma_p/2\hbar k_L = 4$. Several groups have reported lower temperatures in their traps; ultimate temperatures of $3 \mu\text{K}$ are not uncommon [Salomon90; Hamann98]. Typically, polarization gradient cooling works most efficiently in the total absence of magnetic fields. A common strategy is to turn off the anti-Helmholtz coils before performing PGC with the weakened (low intensity and further detuned) optical molasses. Although we can *in principle* turn off the fields for this stage of our experiment, the extremely slow magnetic field ringdown in our interaction chamber (§2.4.2) makes this impractical. While the ultimate temperatures that can be achieved in a MOT are good for many purposes, it is often desirable to have lower temperatures. The most limiting factor in the MOT temperature is reabsorption of scattered photons, a phenomenon that becomes more severe at high densities [Sesko91; Cirac96; Castin98]. When we increase the detuning and decrease the intensity of our molasses during the PGC cooling period we reduce the scattering rate to further cool the sample.

Our laboratory has not made the exploration of cooling techniques a major priority. Instead, we rely upon the results of the many groups that have developed and continue to pursue novel cooling methods. There are several variations on the MOT (with names like grey or dark optical molasses) that have been developed in order to reduce rescattering and have produced temperatures as low as $1 \mu\text{K}$ [Ketterle93; Hemmerich94; Boiron96]. Several promising techniques have been developed in the last few years for cooling atoms (from a MOT) in far-detuned optical lattices. Generally speaking, these techniques work

because confinement in the lattice minimizes heating from rescattering [Cirac96; Wolf00]. One such method is Raman sideband cooling, which has been demonstrated in certain lattice geometries with results close to the recoil temperature [Hamann98; Vuletić98]. One other such method, which we employ, is to trap atoms in a three-dimensional far-detuned optical lattice and cool them with optical molasses.

We have adopted a scheme that is very similar to one developed by David Weiss and his group [Winoto99; Wolf00]. The essential idea is that we load the atoms from our MOT into a three-dimensional, far off-resonance lattice (FORL). The lattice serves to confine the atoms in space while we continue to cool them with a weak optical molasses. Polarization gradient cooling turns out to work (at least) as well in the FORL as it does in free space, and we can achieve very low temperatures while retaining spatial confinement. We then adiabatically release the atoms from the 3D lattice to trade their confinement in position for a further reduction in temperature. The typical temperatures that we can achieve with this technique are lower than $1 \mu\text{K}$. Beyond simply cooling the atoms, we can use this technique to trap and store our atoms while the magnetic field decays in the interaction chamber.

Our lattice configuration is illustrated in Fig. 4.1. Five red-detuned, linearly polarized beams meet at the center of the interaction chamber. Three of the beams are in the plane of the horizontal optical molasses beams, rotated by several degrees about the vertical axis of the chamber from the positions of three of the molasses beams. Two of the beams constitute a retroreflected pair, and the third beam is a traveling wave perpendicular to these beams. These beams are vertically polarized and form a two-dimensional interference pattern in space. The interference pattern contains many local intensity maxima in space. In this particular configuration, the local maxima are stable against small perturbations of the optical phases. We use red-detuned light so that the dipole force attracts atoms to regions of higher intensity, so that the atoms are confined in many small wells. If this lattice configuration is instead realized with blue-detuned light, the intensity maxima form an array of repulsive barriers. We have verified experimentally that in this case the lattice




Figure 4.1: Configuration of the 3D far-detuned optical lattice. The approximate position of the beams is shown with respect to the optical molasses beams (a). Three of the beams are in the same plane as the horizontal molasses beams, and the other two are aligned several degrees from the vertical axis. Each of the lattice beams is linearly polarized, where the polarization vectors are indicated by the small black arrows (b). The vertical beams are detuned by 80 MHz from the horizontal beams.

inhibits motion to some degree, but does not efficiently confine the atomic sample. In the lattice configuration originally described by Weiss's group [DePue99], the horizontal beams are formed by two perpendicular standing waves. An effect of this is that the lattice potential is not stable against small perturbations to the optical phases. To compensate, their retroreflecting mirror locations are stabilized by active interferometric phase locking. There is one possible advantage to working with blue-detuned light, if the lattice were composed of three perpendicular standing waves. In this case, the repulsive force would eventually trap the atoms in the regions of *lowest* intensity. Spontaneous emission from the lattice beams would be minimized in this case, and it should be possible to store the atoms in the dark parts of the lattice for a very long time.

Now let us examine the vertical beams, which are largely independent of the horizontal lattice. The beams are linearly polarized and form a retroreflected pair. The axis of the beam is offset several degrees from the vertical axis so as to not block the vertical MOT beams. This standing wave constitutes a one-dimensional optical lattice and serves to trap

the atoms in the vertical direction, against gravity. The detuning between the horizontal and vertical beams determines the nature of the interference pattern that is formed between the two lattices. The frequency difference between the vertical and horizontal beams is 80 MHz, which is very large compared to the recoil frequency. Therefore, the interference pattern changes much faster than the atoms can respond and so the atomic interactions with the vertical and horizontal beams can be considered to be independent. A secondary effect of the beam independence is that the lattice polarization is effectively linear at every point.

The vertical and horizontal optical lattices provide confinement for our atoms, but they do not (by themselves) provide cooling. Polarization gradient cooling is possible in the FORL so long as the ac Stark shift that the lattice beams induce is nearly independent of the magnetic sublevel m_F [Winoto99]. This condition is satisfied for our configuration because the lattice is linearly polarized and the detuning is large compared to the width of the excited state (hyperfine) manifold. Cooling in the lattice is particularly effective because of several factors. First, PGC works extremely well because we are able to trap the atoms long enough that the magnetic fields in the chamber decay away. Secondly, confinement tends to suppress rescattering of light in the trap, as we have mentioned. We further suppress scattering by reducing the number of atoms that are in the $F_g = 4$ ground state at any given time. When we apply optical molasses light to atoms in the FORL, we turn off the repumping laser. The molasses light at this point is tuned ~ 40 MHz to the red of the $F_g = 4 \rightarrow F_e = 5$ cycling transition. Without any repumping light, nearly all of the atoms end up in the $F_g = 3$ ground state within several milliseconds. Atoms that are in the $F_g = 3$ state do not “see” the molasses light. The low rate of spontaneous scattering events from the optical lattice itself provides enough repumping to occasionally promote an atom to the $F_g = 4$ state so that it will undergo many cooling events before its decay to $F_g = 3$. Rescattering is unlikely because most of the atoms (at any one time) are not participating in the cooling process. (This is much like the behavior in a dark MOT [Ketterle93].) This particular cooling process would be enhanced by a blue-detuned lattice, because the atoms with the largest energies would be those which are most likely to reach regions of higher

intensity. This would selectively repump and cool the atoms with the highest energies.

Beyond simply allowing for good polarization gradient cooling, cooling in the lattice provides one other essential feature: confinement. The temperature that the atoms reach after cooling in the lattice is several microkelvin, comparable to the best results that have been seen in free-space cooling. However, the atoms at this temperature are still localized in their wells. We can now adiabatically lower the well depth by reducing the intensity of the light slowly. In this process, called adiabatic cooling [Chen92], we trade the local confinement in the wells of the 3D lattice for an overall reduction in the temperature of the sample [Kastberg95]. Note that we are only sacrificing *local* confinement—the atomic density averaged over a unit cell of the lattice does not decrease. This stage of the cooling gives us a large improvement over the temperatures that can be achieved with optical molasses alone, and typically brings our atomic sample to below $1 \mu\text{K}$.

The 3D far-detuned lattice beams are derived from the output beam of the Ti:sapphire laser, as shown in Fig. 4.2. The light is picked off from the output beam by an IntraAction 80 MHz fixed frequency AOM that is placed between the laser and the AOM that we use to control the interaction beam. The pick-off AOM is controlled by an IntraAction driver which is in turn controlled by the output of a DS345 arbitrary waveform generator. The light that we have picked off is split by a second IntraAction 80 MHz AOM, which is left on at all times. One third of the light that passes through this AOM is diffracted into the first order output beam which is shifted up in frequency by 80 MHz. The other two thirds of the light passes through undiffracted. The two output beams from this AOM are separated and independently spatially filtered by passing each beam through a $50 \mu\text{m}$ pinhole. The unshifted light is split into two beams of equal intensity with a waveplate and a polarizing beamsplitter cube. These two beams are directed into the chamber to form the horizontal beams. The upshifted light is periscoped up above the chamber to form the vertical lattice beams. Each beam that enters the chamber passes through a final polarizing beamsplitter cube before it enters the chamber to ensure that the polarization is as linear as possible. The typical power of each beam is roughly 90 mW. The three beams are focused onto retrore-

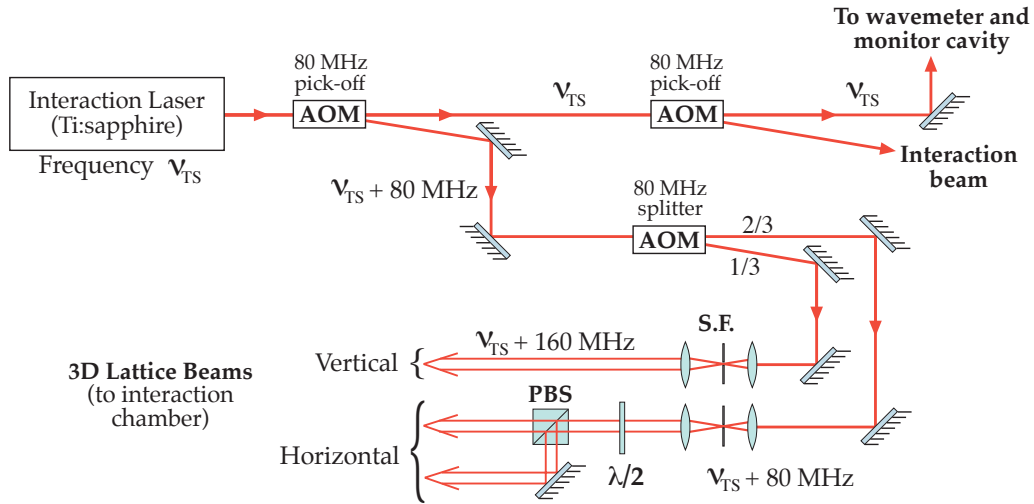


Figure 4.2: Simplified optical setup for the 3D far-detuned lattice. The light is picked off from the Ti:sapphire laser beam (at frequency ν_{TS}) with an 80 MHz acousto-optic modulator (AOM) that is placed between the laser and the AOM that directs light to the interaction beam. When the lattice beam AOM is switched on, it directs light to a second 80 MHz AOM, which splits the light into two components. One third of the light is shifted up in frequency by an additional 80 MHz, and two thirds pass through the second AOM without a frequency shift. The two beams are then separated and passed through independent spatial filters (S.F.). The unshifted beam is split in half with a wave plate ($\lambda/2$) and a polarizing beamsplitter cube (PBS), and these two beams are directed into the interaction chamber as the horizontal lattice beams. The component of the beam that was shifted by an additional 80 MHz is used for the vertical lattice beam pair. In addition to what is shown in this diagram, each beam is weakly focused onto its retroreflector, and passed through a final polarizer before entering the chamber.

flecting mirrors, two on the chamber windows opposite where the horizontal beams enter, and one below the chamber to reflect the vertical beam. In practice, we keep one of the horizontal retroreflectors blocked for our five-beam geometry. The beams are focused onto the retroreflectors so that each beam is the same size ($w_0 = 500 \mu\text{m}$) on its way to and from the retroreflector when it intersects the atomic distribution.

Let us now summarize the experimental sequence that we employ for cooling in the far-detuned 3D lattice. We begin with a sample of atoms that are collected in the MOT, as we have described before. The optical molasses light is at its full intensity and detuned 13 MHz red of the cycling transition. Up until the cooling trigger (§2.7) is issued, the experimental sequence is the same as we have described previously. The first step after the

cooling trigger is to load atoms into the three-dimensional far-detuned lattice. As in the case without the lattices, we detune the molasses light further from resonance to about 37 MHz, and reduce the intensity to 60% of that which we use for the MOT loading phase. The 3D FORL is turned on adiabatically so that it does not heat the sample more than necessary. After 800 μs , the lattice reaches its full intensity. At this point the weak optical molasses, the repump laser and the anti-Helmholtz coils are still on. These components constitute a weak MOT which is left on for a total of 22 ms to help load atoms into the wells of the 3D lattice. After the end of the loading period, we extinguish the repump light and turn off the anti-Helmholtz coils. The optical molasses light is brought back to its full intensity (100%), but the detuning remains at 37 MHz. We now begin a long period of storing atoms in the FORL. The storage period is usually 298 ms, which means that the FORL is on for a total of 320 ms. During the storage period the 3D FORL remains fully bright while the optical molasses intensity is linearly decreased from 100% to 77%. The timescale of the storage period is largely limited by magnetic field decay in the chamber. After this long storage period, the residual magnetic fields in the chamber are no longer a problem for our subsequent state preparation stages. It is also important to note that this timescale is sufficiently long for the polarization gradient cooling to cool the atoms that are trapped in the lattice. The atoms that were not loaded into the lattice initially have plenty of time to gravitationally fall away from the trapped sample. The final stage of the lattice cooling is the adiabatic reduction of the 3D lattice intensity. The intensity as a function of time is lowered according to $I(t) = I_0(1 + t/\tau)^{-2}$ [Kastberg95; DePue99], where our time constant is $\tau = 30 \mu\text{s}$. This function is truncated after 800 μs , at which point the light is already very close to extinguished. When we earlier adiabatically turned on the lattice light, its intensity was increased according to the reverse of this function. The beams which form the optical molasses also constitute a three-dimensional (but near-resonant) optical lattice, and we can slightly decrease the final temperature by adiabatically turning off the molasses beams. These beams are turned down and extinguished with the same temporal profile, but begin shutting off 20 μs before the FORL beams.

Aside from the total length of the storage time, which is determined by magnetic field characteristics, the parameters that we have used are empirically determined by optimizing the final temperature of the atomic sample. Aside from the temperature, it is also important that we retain a reasonably large fraction of those atoms that we initially loaded in the MOT. We can retain between 50% and 90% of the initial atoms for 320 ms in the optical lattice, with 75% a typical value. The most critical parameter in lattice retention is the *alignment* of the 3D lattice beams, not any of the system parameters that we have discussed. There is one additional system parameter that does critically affect the performance of the lattice cooling: the detuning of the lattice beams.

We have found cooling in the lattice to work for detunings as low 12 GHz and at least as high as 70 GHz (to the red of the cycling transition). The performance of the procedure is compromised at low detunings by heating and at high detunings by low retention. The different lattice structure with blue detuning typically results in both poor temperatures and poor retention. We have had the best results for detunings in the range of 25 GHz - 60 GHz. In this range we achieve typical one-dimensional temperatures of 400 nK, which correspond to $\sigma_p/2\hbar k_L = 0.7$. This temperature is measured along the axis of the interaction beam. In the vertical direction, the temperature is slightly warmer, usually about 500 nK, or $\sigma_p/2\hbar k_L = 0.8$. For our experiments, we are primarily concerned with the temperature along the axis of the interaction beam. The temperature along this axis is usually optimized at the expense of temperature in the vertical direction by changing the relative intensities of the 3D lattice beams. The lattice cooling procedure works best, in terms of both retention and temperature, when the lattice beams are made as intense as possible.

Our final atomic distribution after lattice cooling is a mixed state, with most of the atoms in $F_g = 3$, but some in $F_g = 4$. In any of our experiments, it is important to begin with the atoms in only *one* of these ground states. Usually, we optically pump the atoms to the $F_g = 4$ state with a 100 μ s pulse from the repump laser. After this procedure, the 1D temperature is typically near 700 nK. We could also imagine simply removing all of the atoms that were in $F_g = 4$ with a pulse that is resonant on the cycling transition. We

do in fact use a procedure similar to this as a later part of the state-preparation sequence. Beyond this, we also employ a more complex method to pump atoms into the $F_g = 4, m_F = 0$ magnetic sublevel, as we will see later.

4.3 Stimulated Raman Velocity Selection

4.3.1 Overview

As cold as we are able to make atoms with polarization gradient cooling, we are still limited by the recoil temperature. In the next stage of our state-preparation sequence, we use two-photon stimulated Raman transitions to select only a small velocity class of our atomic sample. This technique was first demonstrated by Steven Chu and co-workers [Kasevich91], and we will review some of the essential features here.

We begin by considering the multilevel atom depicted in Fig. 4.3. This atom has two ground states $|g_1\rangle$ and $|g_2\rangle$ separated in energy by $\hbar\omega_{21}$, and a manifold of closely-spaced excited states $|e_n\rangle$. This structure corresponds closely to that of the cesium atom, which has two widely separated hyperfine ground states. Let us suppose that the atom begins in the state $|g_1\rangle$. We now simultaneously expose the atoms to two counterpropagating laser fields




Figure 4.3: This simplified term diagram shows the energy levels that are relevant to stimulated Raman transitions. The two ground hyperfine states $|g_1\rangle$ and $|g_2\rangle$ correspond to the $F_g = 3$ and $F_g = 4$ ground states of the cesium atom. The excited state levels $|e_n\rangle$ are coupled to the ground states by two laser fields of frequency ω_{L1} and ω_{L2} . These laser fields are far detuned from the nearest optical transitions at frequencies ω_{eg1} and ω_{eg2} (Diagram: D. Steck).

of frequency ω_{L1} and ω_{L2} . These beams can induce a stimulated Raman transition to the $|g_2\rangle$ state if the frequency difference between the beams meets a resonance condition. In a simplistic view, the resonance condition for a stationary atom is given by $\omega_{L1} - \omega_{L2} = \omega_{21}$. As before, we operate in a far-detuned regime where the atoms are unlikely to spontaneously scatter photons. In the two-frequency case, this means that there is a large detuning Δ_i ($i = 1, 2$) of each of the laser fields from the single-photon (optical) atomic resonances. When the Raman resonance condition is met, we can view the two laser frequencies as being resonant with a virtual excited state which is far detuned from the excited state manifold. An atom that makes the transition to the other state acquires momentum $2\hbar k_L \equiv \hbar k_{L1} + \hbar k_{L2}$, where k_{Li} is the wave number of light with frequency ω_{Li} . This is to say that besides simply coupling the internal states with the two light fields, we also couple the momentum states of the atom. For a moving atom, the transition is possible if the light fields are doppler shifted such that the resonance condition is met. The momentum spread Δp of the atoms that are transferred to the other state is given by $\Delta p = mc\Delta\omega/(\omega_{L1} + \omega_{L2})$, where $\Delta\omega$ is the linewidth of the Raman transition. The ground state lifetimes are long enough that the linewidth of the Raman transition is in practice limited by the interaction time.

A detailed theoretical consideration of evolution in this system is presented elsewhere [Moler92; Steck01b]. The analysis follows a similar procedure to that in §1.4.1, where the system now reduces to coupled equations of motion for the two ground states. The results take the form of a two-level system, where we can drive oscillations between the states $|p\rangle|g_1\rangle$ and $|p + 2\hbar k_L\rangle|g_2\rangle$. The frequency of oscillation is a Raman Rabi frequency Ω_R that is defined in analogy with the single-photon Rabi frequency Ω . The stimulated Raman resonance condition turns out to be

$$\Delta_R \equiv 4\omega_r \left(\frac{p + \hbar k_L}{\hbar k_L} \right) + (\Delta_{L2} - \Delta_{L1}) + (\omega_{AC2} - \omega_{AC1}) = 0 \quad , \quad (4.1)$$

where ω_{AC1} and ω_{AC2} are the ac Stark shifts of the two ground state levels. The overall resonance condition is the sum of three parts: the doppler shifts, the simplistic resonance condition $\omega_{L1} - \omega_{L2} = \omega_{21}$, and the ac Stark shifts of the atom in the two laser fields. This

model does not account for changes in the value of ω_{21} that can be induced by *external* fields. Zeeman shifts caused by the residual magnetic field are a potential problem for our experiments for this reason.

By exposing the atoms to the Raman laser fields, we can drive Rabi oscillations for the entire class of atoms that meet the resonance condition $\Delta_{\text{R}} = 0$ (4.1). Assuming that all of the atoms begin in $|p\rangle|g_1\rangle$, the population ρ of the state $|p + 2\hbar k_L\rangle|g_2\rangle$ is given by $\rho(t) = \sin^2(\frac{1}{2}\Omega_{\text{R}}t)$. Usually we wish to drive a “ π -pulse,” by leaving the beams on for a period $t_{\pi} = \pi/\Omega_{\text{R}}$. In this case, an atom in $|p\rangle|g_1\rangle$ makes a transition to $|p + 2\hbar k_L\rangle|g_2\rangle$ (or vice versa) with unity probability. The strenuous condition $\Delta_{\text{R}} = 0$ is not met by most of the atoms and the population is instead given by a generalized Rabi formula,

$$\rho(t) = \frac{\Omega_{\text{R}}^2}{\Omega_{\text{R}}^2 + \Delta_{\text{R}}^2} \sin^2\left(\frac{1}{2}\sqrt{(\Omega_{\text{R}}^2 + \Delta_{\text{R}}^2)} t\right) . \quad (4.2)$$

This equation implicitly assumes that Ω_{R} is constant over the time interval $(0, t)$, which is the case for a square pulse. So, for a square pulse we see that the probability of making the transition oscillates as a function of Δ_{R} . If the pulse time is $t = t_{\pi}$, then the function is a maximum at $\Delta_{\text{R}} = 0$, surrounded by broad oscillatory lobes that reach zero periodically. The width of this distribution is determined by the interaction time, and varies as $1/t$, where t is the total time of the π pulse. With a longer pulse, the energy of the selected velocity group is more clearly defined, and the size of the oscillatory structure is decreased. The effective Raman detuning Δ_{R} depends upon the velocity of each atom, so when we apply a Raman pulse to the distribution it actually excites atoms within a small range of momentum. To fully account for the time evolution of the excited state, we must consider the effect of Eq. 4.2 on the entire atomic momentum. The result is that the excited state population does not undergo perfect sinusoidal oscillation, but rather exhibits damped oscillation. It is possible to reduce the side lobes by choosing a non-square pulse shape, as has been demonstrated by Chu and co-workers [Kasevich92]. If the intensity of the pulse is chosen with a Blackman pulse envelope, the frequency spectrum of the pulse has only very low power components outside the central region. Despite this advantage, we decided to continue to

use square pulses in our experiment. The chief motivation for this decision is that the ac Stark shifts change as the intensity of the pulse is modified. In order to keep a set of atoms in resonance with the Raman transition, the frequency difference between the beams would have to be modified as the intensity was changed. In our setup we use radio-frequency (RF) electronics to drive the AOMs that control the stimulated Raman laser beams. These electronics cannot perform this real-time, arbitrary frequency modulation while maintaining high enough stability for our experiments.

Let us now describe the basic procedure that we use in our experiment to perform velocity selection. We begin with a cold sample of atoms in the $F_g = 4, m_F = 0$ atomic state, obtained after cooling in the 3D lattice and optical pumping. We then use a stimulated-Raman “tagging” pulse to promote only those atoms within a certain velocity range to the $F_g = 3, m_F = 0$ state. Finally, we apply resonant light on the $F_g = 4 \rightarrow F_g = 5$ line to push away all the atoms that remain in the $F_g = 4$ state. We now have a sample of atoms in the $F_g = 3$ state with a narrow velocity distribution, which is often a desirable initial condition for an experiment. We use a transition between two $m_F = 0$ states because splitting between these states is highly insensitive to Zeeman shifts. We have skipped over some of the details here, which we will come to in due course.

One can imagine extending this state-selection procedure into a true *cooling* process by selecting the atoms near zero velocity, recycling the non-selected atoms into a new thermal distribution, and repeating. Stimulated Raman cooling based upon this principle has been demonstrated [Kasevich92]. More technically, the final distribution should have zero velocity, so the momentum class that is initially tagged must be at $\pm 2\hbar k_L$. In the process that we have described so far, we can tag atoms from the $|p = 2\hbar k_L\rangle|F_g = 4\rangle$ state into the $|p = 0\rangle|F_g = 3\rangle$ state. We can reverse the direction of momentum transfer by reversing the sign of $\omega_{L1} - \omega_{L2}$, so that we couple the states $|p = -2\hbar k_L\rangle|F_g = 4\rangle$ and $|p = 0\rangle|F_g = 3\rangle$. We can then, by changing the frequency of (say) ω_{L1} , drive π -pulses that move atoms from both positive and negative momentum classes towards zero. The recycling procedure would have to involve the use of resonant light to redistribute the velocity distribution of the $F_g = 4$

atoms. Implementation of stimulated Raman cooling is made difficult by several factors. The most important of these is the extreme sensitivity that the system has to the effects of external magnetic fields. The residual magnetic fields in our interaction chamber are significant enough (typically at least 20 mG) that we do not expect Raman cooling to work, at least in its standard implementations. Originally, we planned to perform a slightly modified version of Raman cooling in our experiment, and our optical setup has the capability of switching the detuning of the light. Once we had the capability of velocity-selective tagging, we briefly attempted several variations on the cooling technique. Although we did not pursue any of these tests in detail (in the interest of time), we were not able to exceed the performance of simple tagging alone. The chief disadvantage of performing only the single velocity selective tagging is that we end up “throwing away” most of our distribution. In the chaos-assisted tunneling experiments, we usually ended up performing the experiments with about 0.1% of the atoms that were initially gathered in the MOT. With such a small atomic distribution to work with, we often needed to average experiments together to increase the signal to noise ratio.

4.3.2 Experimental setup for velocity selection

We now come to the experimental setup that is used to generate the Raman laser fields. A simplified overview of the optical components of this system is shown in Fig. 4.4. The Raman beams are derived from the same far-detuned Ti:sapphire laser that we use to form the interaction beam and the 3D far-detuned lattice. The output beam of the Ti:sapphire laser is multiplexed by a series of acousto-optic modulators. We have already seen (Fig. 4.2) that the beam passes through the 3D lattice control AOM and the interaction beam control AOM. An IntraAction 40 MHz AOM follows these two, and is used to control the intensity of the Raman beams. The initial frequency of the Raman light, the *carrier frequency* ν_c is a fixed 40 MHz above the Ti:sapphire laser frequency, $\nu_c = \nu_{\text{TS}} + 40$ MHz. The Ti:sapphire laser is kept at a known frequency ν_{TS} by using the wavemeter and monitor cavity described in §2.3.2.

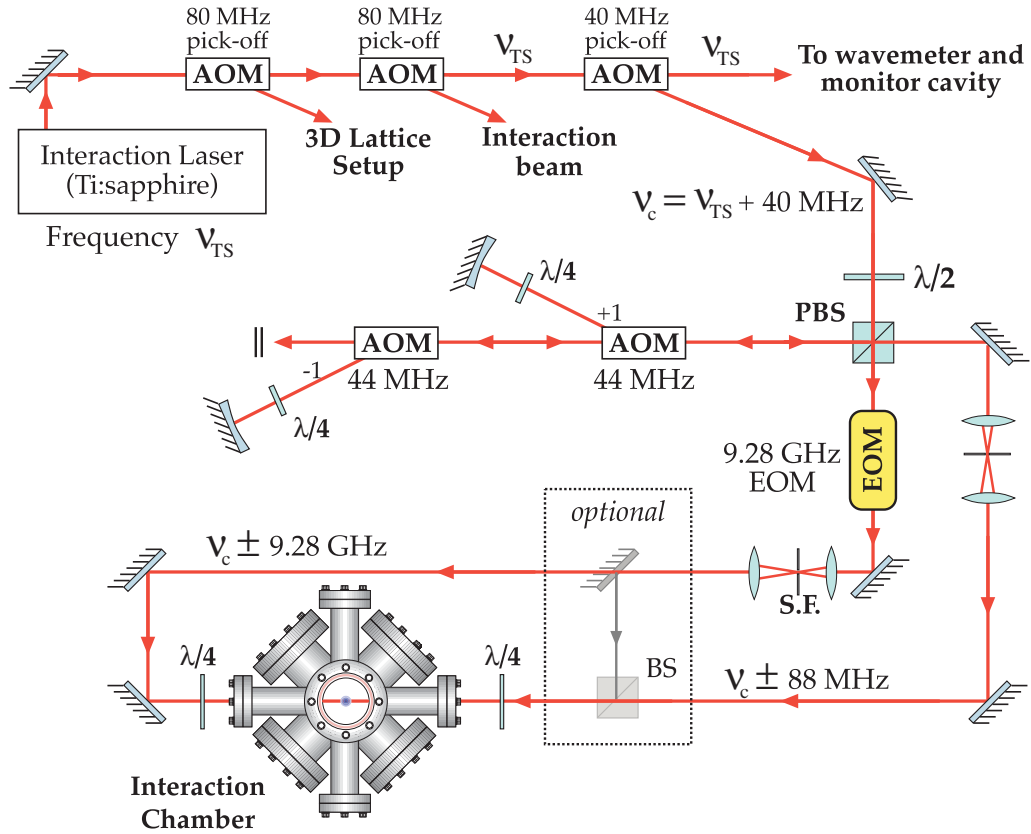


Figure 4.4: Simplified optical setup for stimulated Raman transitions. The light is picked off from the Ti:sapphire laser beam (at frequency ν_{TS} with a 40 MHz AOM that is placed between the AOM that directs light to the interaction beam and the wavemeter. When the pick-off AOM is switched on, it directs light at the carrier frequency $\nu_c = \nu_{TS} + 40$ MHz to a variable beamsplitter. Half of the light passes through a 9.28 GHz EOM and is spatially filtered before entering the chamber. The other half of the light is double passed through one of two tunable 44 MHz AOMs before being spatially filtered and entering the chamber opposite the beam from the EOM. Both beams pass through quarter-wave plates ($\lambda/4$) just before the chamber so that the beams become circularly polarized. In an optional setup selected by the use of kinematic stages (indicated by a dotted box) the light from the EOM arm of the setup is coupled with the light from the AOM arm with a non-polarizing beamsplitter cube (BS).

The light for the Raman beams is first split into two equal components with a variable beamsplitter consisting of a half-wave plate ($\lambda/2$) and a polarizing beamsplitter cube. Each half begins with a total power in the neighborhood of 200 mW. The first half of the light travels through a New Focus model 4851 9.28 GHz electro-optic phase modulator (EOM). This modulator is driven by a stable 9.28 GHz source that we will describe later, and

serves to put sidebands at ± 9.28 GHz on the beam. We measured the degree of modulation by looking at this spectrum with a monitor cavity identical to the one that we use to monitor the main Ti:sapphire beam. In this measurement, we found that the EOM converts roughly 7% of the carrier into each sideband. The aperture on the EOM is very small (1×2 mm), so the beam is weakly focused through the EOM with a 400 mm (focal length) lens. An identical lens after the EOM recollimates the beam, and it is then spatially filtered with a $40 \mu\text{m}$ pinhole. The EOM transmits roughly 90% of the light that enters it— the losses are presumably due to clipping at the input and output apertures. This clipping also distorts the beam somewhat and after spatially filtering, we typically have an output power near 150 mW.

The other half of the Raman light, which does not pass through the EOM, is directed to a matched pair of tunable 40 MHz acousto-optic modulators from IntraAction. The AOMs are driven near 43.5 MHz by a stable RF setup that we will describe later. They are arranged in series as shown in Fig. 4.4. Both AOMs are double-passed, the first on the +1 order, and the second on the -1 order. Only one of the AOMs is on at a given time, so that the output frequency of the double-pass setup is at the single frequency $\nu_c \pm 87$ MHz. To improve the maximum efficiency of the setup, the beams are weakly focused as they pass through the AOMs, initially with a 400 mm lens. Both retroreflecting mirrors are 99% reflectivity spherical mirrors with a 50 cm radius of curvature, and are mounted on 1D translation stages to allow adjustment of the focus. The two double-pass setups use the same polarizing beamsplitter cube, and the output of the setups is finally spatially filtered with a $35 \mu\text{m}$ pinhole. In all of the experiments described in this dissertation, we used only the double-pass setup that shifts *up* the frequency of the light. Typically, the double-pass setup is only about 45% efficient, and we end up with 75 mW of light after the spatial filter.

After the spatial filters, the beams from the EOM and AOM arms of the optical setup enter the chamber from windows opposite each other and are made to counterpropagate by carefully aligning the beam paths. The beams both have a waist of about 2 mm, and they are aligned very close to the path of the standing wave that provides the interaction

potential. The angle between the Raman beams and the interaction beam is about 1° in the horizontal plane. The small angle between the beams ensures that our velocity selection is with respect to the axis that we actually care about, that of the interaction beam. The Raman beams need to be circularly-polarized for the $m_F = 0 \rightarrow m_F = 0$ transition, so the two Raman beams each pass through a final polarizer and a zero-order quarter-wave plate just before each beam enters the interaction chamber.

In the stimulated Raman transitions that we have described so far, the beams are counterpropagating and this is the source of the velocity dependence in the Raman resonance condition. If we instead switch to a copropagating scheme, the resonance condition can in principle be met simultaneously for atoms in every velocity class. In this case, the resonance condition is met if the proper frequency difference between the beams is chosen to account for any level shifts. In our optical setup, we can insert a mirror and beamsplitter to overlap the light from the EOM with the light from the AOM. In this case, the beams with two frequency components reach the atoms from the same side. Since we are then insensitive to the velocity class, this configuration makes it much easier for us to evaluate the ac Stark shift and shifts due to residual magnetic fields.

We have introduced the means by which we generate some of the frequencies on the




Figure 4.5: Levels and frequencies for stimulated Raman tagging. The cases shown are where the AOM double pass setups use the +1 order (a) and (b) where they use the -1 order (Diagram: D. Steck).

Raman beams. Let us now relate those frequencies to the energy levels in the cesium atom. An energy level diagram with the experimental frequency values is shown in Fig. 4.5. The beam which passes through the EOM has frequency components at ν_c and $\nu_c \pm 9.28$ GHz. The beam which is derived from the AOM arm of the setup is entirely at either $\nu_c + 87$ MHz or $\nu_c - 87$ MHz. The ground state hyperfine splitting $\hbar\omega_{21} = (E(F_g = 4) - E(F_g = 3))$ of the cesium atom gives the frequency ω_{21} of the clock transition, with the well known value $2\pi \cdot 9.192\,631\,770$ GHz. The difference between this value and the 9.28 GHz drive is 87.368 230 MHz, which we obtain by driving the AOMs at 43.684 115 MHz. The actual resonance condition is of course subject to doppler, ac Stark, and Zeeman shifts. For these reasons, we actually send a signal at the modified frequency $43.684\,115$ MHz $+\delta$, and the light is shifted by $87.368\,230$ MHz $+2\delta$. We can also use the sideband at $\nu_c - 9.28$ GHz to achieve the resonance condition, using the other double-pass setup, with $\nu_c - 87.368\,230$ MHz. These two realizations reverse the sign of the frequency difference, and therefore the direction of the momentum kick that is associated with the Raman transition. Although we only use one sideband of the light from the EOM, the other sideband and the carrier frequency are not important (except for the possibility of spontaneous emission). The carrier does form a standing wave with the light from the AOM at $\nu_c \pm 87$ MHz. Note however, that 87 MHz is extremely large compared to the recoil frequency, and this standing wave moves at a high enough velocity that it can safely be ignored.

The stability of the frequency difference between the relevant components of the stimulated Raman beams is critical for several reasons. First off, in Eq. 4.1 we see that if the frequency difference between the two beams is $\Delta_R = 4\omega_r$, then we are resonant with a momentum class moving at $\hbar k_L$. Recall that the recoil frequency is $\omega_r = 2\pi \cdot 2.0663$ kHz. This means that if we wish to select a given velocity group, we must have stability much better than several kilohertz. A much more stringent requirement is imposed by the narrow width of the velocity group that we wish to select. We are interested in selecting a subrecoil sample of atoms— one with a width of as low as $\sim \hbar k_L/20$. In order to do this, we need to have frequency stability much better than $\omega_r/20 \approx 100$ Hz. This stability must

be maintained over the total length of time that is required for velocity selection. We can estimate the π -pulse time that will produce this velocity selection by setting $\Delta E \Delta t = \hbar/2$, with $\Delta E = \hbar\omega_r/20$. The pulse time is then $\Delta t = 10/\omega_r \sim 800\mu s$. The actual width selected by an $800\mu s$ pulse is a slightly larger $0.06\hbar k_L$, because of the shape of the distribution selected by a square pulse.

All of the laser light that constitutes the stimulated Raman beams is derived from the same Ti:sapphire laser. This is important because short-term fluctuations in the laser frequency are common to each frequency component of the light and therefore cancel out in the difference between any two. Beyond this, it is important to make sure that the signals used to drive the EOM and the AOMs are very stable. The timebase of our signal generators are all referenced to the same EFRATOM rubidium clock that we use to stabilize our other instruments. The rubidium clock outputs a stable 10 MHz reference signal.

The 9.28 GHz signal for the EOM is generated by a Delphi Components Inc. dielectric resonant oscillator (DRO). This DRO takes as its input a 40 MHz signal (obtained by doubling the 10 MHz reference signal twice), and produces an output signal that is an exact frequency multiple of the input. The 9.28 GHz signal is isolated by Sierra Microwave Technology model SMC-8010 circulator which ensures that any backreflection is directed to a $50\ \Omega$ power terminator rather than into the DRO output. After the circulator, the 9.28 GHz signal is amplified with a QuinStar Technology Inc. model CPA09092535-1 solid state amplifier. This amplifier is specified to give 25 dB of gain, up to a maximum of 35 dBm, at 9.28 GHz, although its performance may actually be slightly below this level. The amplifier dissipates a large amount of heat in operation and is sensitive to temperature. For these reasons, it is mounted to a massive heat sink from an old oscilloscope, and the heat sink is cooled with the assistance of two standard computer fans. The high-power output signal from the amplifier is transmitted to the EOM through a 1 m SMA cable constructed of Times Microwave Systems LMR-400 cable. LMR-400 is a flexible cable in the RG-8 form factor and has relatively low loss in our frequency range.

The signal to drive the AOMs is generated by a slightly more complex setup. Beyond

simply needing stability, we also need to be able to tune the output frequency. We begin with a 150 MHz signal generated by one of two synthesized waveform generators, a Fluke model 6080A/AN synthesizer or an Hewlett-Packard model 8662A synthesizer. These can be used with their FM analog inputs to sweep the output frequency over some small range. For the HP, the range is ± 25 kHz, and the Fluke can sweep over ± 1 MHz. The wide sweep range of the Fluke has made it very useful for both looking for the first signal after changes to our system and for certain types of calibrations. The Fluke has the significant problem that its output frequency tends to slowly drift as a function of time in a not-obviously-bounded manner, and we have observed deviations on the order of 1 kHz. While this is not a problem for sweeps on the order of 1 MHz, it does prohibit using the synthesizer for any “important” applications, including velocity selection. For more critical applications that do not require wide sweeping ability, we have used the HP synthesizer. When the HP’s FM input was enabled, it too has the problem of slow drifts. However these drifts are of order 100 Hz, and do not pose as much of a problem for us. For our experiments that involve velocity selection, we avoid the drifts entirely by turning off the FM input and letting the computer directly program the desired output frequency over the GPIB interface. In this case, the stability of the HP synthesizer’s output is stable to within our ability to measure it, less than 1 Hz/day. The signal is usually not at *exactly* 150 MHz— it is this signal that we change in order to change the velocity class that we interact with. We double this signal to obtain a signal that is close to 300 MHz. The original purpose for this doubling was to increase the effective range of our FM sweeps to be as high as ± 2 MHz, although we rarely actually need this. This 300 MHz signal is then mixed with a signal near 343.7 MHz and low-pass filtered to obtain an output signal near 43.7 MHz. The 343.7 MHz signal is produced by a WaveTek model 2047 synthesizer and is kept at a constant frequency during data runs. The exact value is sometimes changed to recenter the stimulated Raman resonance condition within the sweep range of the 150 MHz signal— a change that corresponds to compensating for the ac Stark shift or residual net motion of the atomic sample. The 43.7 MHz signal is amplified and sent to one of two IntraAction model PA-4 power amplifiers. These amplifiers drive the

two double-passed acousto-optic modulators in the stimulated Raman setup. Depending on which of the two amplifiers has the input signal at a given time, the AOM arm of our setup will shift the frequency of the light either up or down by 87 MHz.

4.3.3 Optical pumping and pushing

The state-selection procedure that we have described requires the use of resonant light to select the initial state and to remove atoms that have not been tagged. As previously described, we begin with lattice-cooled atoms in a mixture of the $F_g = 3, 4$ ground states. The most simple form of the optical pumping is the application of repumping light alone. This transfers the entire population to the $F_g = 4$ state, in a mixture of magnetic sublevels. For our velocity selection to be insensitive to magnetic fields, we rely upon a stimulated Raman transition between the $F_g = 4, m_F = 0$ and $F_g = 3, m_F = 0$ states. This procedure greatly reduces the complexity of the problem, as we specialize to a single transition out of the 15 possible.

Let us now summarize the optical pumping procedure that we use to accumulate atoms in the $F_g = 4, m_F = 0$ state. We expose the atoms to linearly polarized light that is resonant with the cesium $F_g = 4 \rightarrow F_e = 4$ transition. The dipole matrix element for the $F_g = 4, m_F = 0 \rightarrow F_e = 4, m_F = 0$ transition is zero. The $F_g = 4, m_F = 0$ state is then a “dark” state in that atoms in this state are not affected by the resonant light. Atoms that did not begin in this state then accumulate there after several spontaneous emission cycles. This pumping light enters the interaction chamber from the top as an independent beam. We prevent atoms from accumulating in the $F_g = 3$ state by turning the repumping laser light on during the optical pumping. Besides this, we use a magnetic bias field of 1.5 G along the polarization axis of the optical pumping beam, which is along the axis of the interaction beam. The bias field is much stronger than any other residual and possibly fluctuating magnetic fields that may be present in the interaction region, and defines a clear quantization axis for optical pumping and stimulated Raman processes. We realize the bias field by turning on an increased current in two of the Helmholtz coils that are usually used

to null out ambient fields. The optical pumping sequence can be performed after the atoms are released from the 3D lattice, but we have found it equally acceptable to perform the procedure during the last moments that the 3D lattice is still on. The 3D lattice intensity is gradually decreased and is truncated after 800 μs . Typically, the optical pumping light is turned on 66 μs before the end of this shutoff period, and is left on for 50 μs . The repumping light is turned on at the same time, but stays on until the 3D lattice truncation, because it is extremely important that no atoms remain in the $F_g = 3$ state after the end of this stage. The magnetic bias field is turned on 200 μs before the optical pumping light so that the field has time to ramp up to its full value and stabilize. We can pump at least 95% of our lattice-cooled atoms into the $F_g = 4, m_F = 0$ state, but we are using a process based upon spontaneous emission, and there is significant heating associated with this process. After full optical pumping, the temperature of the distribution is about 3 μK ($\sigma_p/2\hbar k_L = 1.9$). In our tagging process which follows the pumping, we only select very cold atoms that are in the $F_g = 4, m_F = 0$ state. The optical pumping process serves to increase the number of atoms that we tag by collecting atoms in all 17 ground states of the cesium atom into the state that we select from. However, the heating associated with our pumping procedure reduces the number of atoms that we can tag because it reduces the number of atoms at the low velocities that we wish to tag. As it turns out, we gain much more from our pumping than we lose to heating, and pumping to $F_g = 4, m_F = 0$ significantly increases our brightness relative to operating with repump-only optical pumping.

One of our most important calibration procedures is the alignment of the magnetic fields that we use in our experiment. There are two parts to this, nulling out the ambient fields, and aligning the bias field that we use for optical pumping. For the case of nulling out the magnetic fields, we perform a series of individual experiments and attempt to minimize the total width of the Zeeman-split Raman spectrum. The procedure is very similar to how we perform our usual experiments using stimulated-Raman transitions. In each experiment, we pump the atoms into the $F_g = 4$ states with repump-only optical pumping. We then tag a set of atoms to the $F_g = 3$ state, push away the remaining $F_g = 4$ atoms, and count the total

number of atoms by imaging with the CCD and adding up the total fluorescence. We repeat this procedure many times (typically 50 per scan) with a new Raman detuning frequency at each point in the scan. After the scan is complete, we see a peak in the spectrum for each of the Zeeman-split Raman transitions. We then make a minor manual adjustment to the settings on the Helmholtz coil current controller and repeat the scan. We continue this procedure until the total width of the Zeeman-broadened spectrum indicates that the residual fields are of order 20 mG, which is about as well as we can cancel the fields out. The fields tend to drift slowly as a function of time, and we generally do not need the field to be much better than 100 mG, so long as we are using the bias field in our experiments. Frequently, we conduct this type of measurement with wide sweeps of the detuning, for which we use the Fluke synthesizer in the drive electronics for the Raman AOMs. To assist the accuracy of the coarse sweeps, we commonly chirp the Raman detuning over a frequency interval equivalent to the frequency step size. Furthermore, we commonly use the copropagating setup for these calibrations so that we interact with atoms in a wider velocity range. The procedure for aligning the bias field is very similar to the one that we use to zero the fields. We again scan the frequency difference of the Raman beams, but we include the full optical pumping sequence in each experiment, rather than just the repump-only optical pumping. This involves turning on the magnetic bias field, and leaving it on while we tag the atoms to $F_g = 3$. With the bias field on, each of the Zeeman sublevels is clearly separated in the scan. Using this procedure, we can determine the relative populations of the $F_g = 4$ magnetic sublevels. We then adjust the bias field orientation by optimizing the fraction of atoms that are pumped into the $F_g = 4, m_F = 0$ state.

We now come to the process of removing the unwanted $F_g = 4$ atoms after tagging. We use an independent beam that is resonant with the cesium $F_g = 4 \rightarrow F_e = 5$ cycling transition to push away these atoms. This beam enters the interaction chamber with the Raman beam that travels through the 9.28 GHz EOM. The beams are combined with a polarizing beamsplitter close to the chamber, and both beams pass through the quarter-wave plate before entering the chamber. The pushing beam is then circularly polarized and

travels along the quantization axis of the magnetic field. The pushing beam passes through a 25.4 mm focal length lens just before entering the chamber, and rapidly diverges while it is in the chamber, travelling another 50 cm or so before actually hitting the atoms. This ensures that the beam covers all of the atoms uniformly. Besides this, we need very little of this light to perform the pushing. The light in a cycling transition is particularly efficient because it is circularly polarized along the quantization axis: this light only couples m_F to $m'_F = m_F + 1$. We then expect that the atoms eventually fall into the $F_g = 4, m_F = 4 \rightarrow F_e = 5, m_F = 5$ cycling transition, from which there is no (dipole allowed) escape. We use a very-low intensity pushing pulse to minimize the probability of excitation to one any of the other excited states. To perform the pushing itself, we turn on the pushing light at an very low intensity for a fairly long time, typically 800 μs . We calibrate the pushing pulse by altering the intensity and the duration of the pulse while monitoring the atoms that are left behind. We can accelerate nearly the entire set of $F_g = 4$ atoms to a momentum of over $100 \cdot 2\hbar k_L$, which is outside the region that we detect with our ballistic-expansion technique. By using a long CCD exposure time of 100 ms we can detect that about 0.3% of the atoms are typically left behind. These atoms form a broad background in the momentum distribution that is usually ignorable.

The optical setup for the pumping and pushing beams is shown in Fig. 4.6. These two beams are created by manipulating the output light of the DBR diode laser, which is otherwise used to generate the optical molasses light for the MOT. The output of this laser is kept at 195 MHz below the $F_g = 4 \rightarrow F_e = 5$ cycling transition. The light for the optical molasses is switched and tuned by a double-passed tunable 80 – 100 MHz AOM. The zero-order light that is not deflected by this AOM is passed to an AOM that operates at the fixed frequency of 56 MHz. This AOM downshifts the frequency of the beam to 251 MHz below the cycling transition, which places the output on resonance with the $F_g = 4 \rightarrow F_e = 4$ transition (See Fig. 2.2). This is the light that we use for optical pumping to the $F_g = 4, m_F = 0$ state. Finally, the zero-order output from this AOM is double-passed through a fixed-frequency 97.5 MHz AOM which produces output at the frequency of the

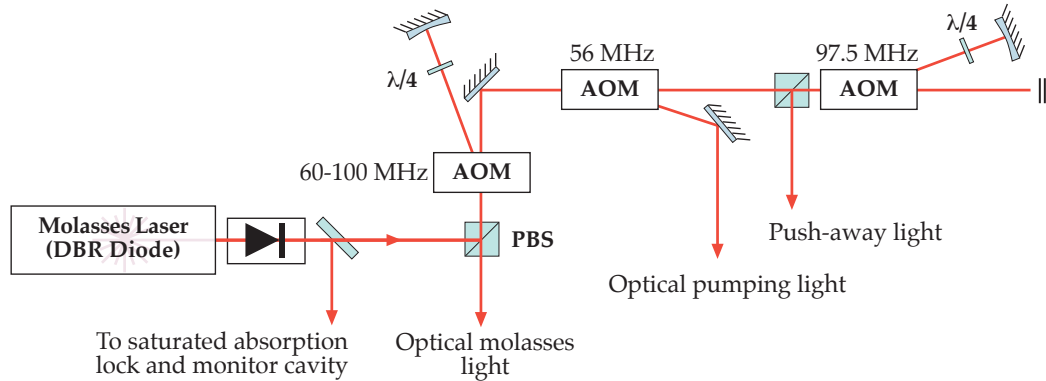


Figure 4.6: Optical setup for the optical pumping and push-away beams. These beams are derived from the DBR diode laser that provides the optical molasses light. The zero-order output of the double-passed optical molasses AOM is directed first to a 56 MHz AOM that is used for optical pumping to the $F_g = 4, m_F = 0$ state. The zero-order output of this AOM is double-passed through an AOM at 97.5 MHz to provide the light that is used to push away $F_g = 4$ atoms in the state selection process.

cycling transition. This beam is used to push away the remaining $F_g = 4$ atoms at the appropriate stage of the state preparation process. Both the pumping and pushing beams are spatially filtered (not shown) before being directed to the chamber.

4.4 State Preparation by Atomic Sliding Motion

The lattice cooling and subsequent stimulated Raman velocity selection allow us to prepare atoms within a narrow range of momentum. However, the atoms that we have tagged are nearly uniform in position, over the scale of one well of our interaction beam. If we wish to study local structures in phase space, we would like to begin with a distribution that is localized (compact) in both position and momentum. The relevant phase space is realized by modulating the same standing wave that we used in the kicked rotor experiments. The exact form of the modulation is not terribly important at the moment, it is enough to note that the spatial periodicity of the potential is again half the wavelength of the light. Besides localizing a sample in phase space, we would like to be able to translate the sample to center it at places in phase space other than $x = 0, p = 0$.

We have developed a series of manipulations that we can use to achieve these goals,




Figure 4.7: Illustration of the state preparation by atomic sliding motion (SPASM) sequence, showing the localization and translation of an atomic sample in a representative $x - p$ phase space (Diagram: D. Steck).

illustrated in Fig. 4.7. This procedure is a modification of an earlier conception that we described in [Klappauf99]. The initial condition for this sequence is the result of the Raman




Figure 4.8: Schematic representation of the experimental sequence for experiments with state preparation. This diagram shows the major steps in the experiment and the associated optical processes. Note that the repumping beam is not depicted in this figure, although it is on during the MOT loading, optical pumping, and imaging stages.

velocity selection, with subrecoil momentum width and uniform in position across the phase space. We then load the atoms into the same 1D standing wave that we use to form the interaction potential. We turn on this light on adiabatically, using the same functional form and time constant that we use to adiabatically deepen the 3D FORL. In the interest of minimizing the total time of our state preparation sequence, the total time of deepening the 1D lattice is $300 \mu\text{s}$ instead of $800 \mu\text{s}$. As the well is deepened, the atoms localize spatially

in the bottom of the wells. This localization can only occur at the expense of the spread in momentum, which can only change in units of $2\hbar k_L$. Because the initial momentum spread of the distribution is much smaller than this interval, these discrete momentum steps are visible as stripes in the momentum in Fig. 4.7. We can gather some idea of the dynamics from considering the band structure of the optical lattice. The initial atomic distribution in p is much smaller than the unit cell of our lattice in momentum space, which spans $(-\hbar k_L, \hbar k_L)$. (We can see the size of the unit cell intuitively by noting that a first-order Bragg transition couples $p = \pm\hbar k_L$.) The narrow momentum distribution is thus adiabatically loaded into the lowest energy band of the optical lattice. For deep potentials with many bound states (like the ones that we use), the lowest band closely approximates the harmonic oscillator ground state. We then expect that the distribution in the standing wave wells is very close to the ground state of the harmonic oscillator, which is a minimum-uncertainty Gaussian wave packet. More technically, this argument only applies to the *outline* of the wavepacket. Within this outline, there is additional confinement into the narrow slices that we have described. The extremely narrow width of these slices in momentum implies that there is a high degree of uncertainty in position, spread over many wells of the standing wave.

The procedure that we have described thus far has produced a near-minimum uncertainty wave packet centered at the bottom of the well in position and centered at zero in momentum. Our next task is to translate this wave packet to other locations in the phase space. We begin the next stage by rapidly translating the phase of the 1D lattice by about one quarter of a lattice period. (We will describe the phase control system shortly, in §4.4.1.) Since the potential is shifted, the free evolution of the atomic distribution lets it “slide” down the potential towards zero, acquiring momentum as it does so. This process is well described by a macroscopic analogy: imagine that we place a stone in the bottom of a smooth bowl. If we suddenly shift the position of the bowl, the stone will slide down towards the center of the bowl. In the experiment, we can jerk the phase of the light and then watch the atoms oscillate back and forth in the potential. At the moment when the atoms reach their maximum velocity, we know that they have returned to the minimum of the potential

Timing Diagram: Experiments w/ Raman Tagging




Figure 4.9: Simplified timing diagram for experiments involving state preparation.

at $x = 0$. If we turn off the potential at this point, we succeed in translating the initial wavepacket in momentum. The total boost in momentum that we achieve by this procedure is determined primarily by the well depth of the lattice during this period. Since we have already placed the distribution at some nonzero momentum, it is also straightforward to place the atomic distribution at any *position* in phase space. In this case, we simply allow a free drift period to relocate the atomic distribution in position. A slight disadvantage of this procedure is that the individual momentum slices within the wavepacket move at different velocities and tend to separate during the drift time. A second complication with these procedures is that the sinusoidal wells of our standing wave introduce anharmonicity to the problem. However, simulations of the state preparation procedure suggest that the distortion introduced into the outline of the final states is not severe. Despite the distortion, the internal “ladder” of momentum states (that we initially selected with the Raman tagging) is preserved because the interactions with the 1D lattice are coherent dipole force interactions.

We have now described a procedure for localizing the atomic position in both position and momentum, and subsequently centering that distribution at arbitrary locations in phase space. We have named this procedure “State Preparation by Atomic Sliding Motion,” or “SPASM” for short. In our data runs, we begin the experiment itself immediately after the SPASM sequence, usually simply by changing the type of modulation that is applied to the standing wave.

The SPASM sequence itself is only a means of preparing an initial state in phase space. In order to be able to apply this method, we must begin with an extremely cold sample which is in our case generated by the lattice cooling and stimulated Raman velocity selection. After the state preparation is complete, we create the interaction potential and detect the final momentum distribution by ballistic expansion as we have described previously. The overall sequence of these experiments is illustrated in Fig. 4.8. Naturally, the computer-control sequence for this sequence is much more complicated than it was for the kicked rotor experiments. A timing diagram for experiments with the state-preparation sequence

is shown in Fig. 4.9.

4.4.1 Optical phase control

In the experiments with the kicked rotor that we have described previously, the phase of the interaction beam was determined by the position of the retroreflecting mirror. This mirror was rigidly fixed to the vacuum chamber window flange for stability. In the experiments utilizing the SPASM sequence and in the squeezing experiments that we will come to later (Chapter 6), we need to be able to control the phase of the interaction beam. We control the phase of the light by placing an electro-optic phase modulator between the interaction chamber and the retroreflecting mirror.

The optical setup for the phase control is shown in Fig. 4.10. We have removed the retroreflector that was previously attached to the interaction chamber and added a vibration-isolated optical breadboard platform near the edge of the interaction chamber. The platform is an irregularly shaped sandwich of two layers of 1/2" aluminum jig plate around a layer of 1/2" soft Sorbothane, a viscoelastic damping material. The top surface of the platform is about 2" below the interaction beam axis. The interaction beam crosses through the chamber before crossing to this platform. On the platform, it is weakly focused by a 300 mm focal length lens onto a retroreflecting mirror. Because the lens is positioned so as to focus the light onto the mirror itself, the lens serves to recollimate the beam before it reenters the interaction chamber. Between the lens and the mirror is a Conoptics, Inc. model 360-40 lithium tantalate electro-optic phase modulator. The lens serves to reduce the size of the beam so that it is not clipped while it passes through the EOM, which has a 2.7 mm (square) clear aperture. The EOM is nearly coaxial with the interaction beam, however it is tilted off of the beam axis by a couple of degrees so that any reflection off of the EOM surfaces is not in the direction of our atomic sample.

The phase modulator is driven by a Conoptics model 302 driver, which provides a bias of ± 400 V. We have measured the phase of our standing wave to be shifted by a full period when the EOM control voltage is shifted by 250 V. Although we can in principle




Figure 4.10: Phase control setup for the interaction beam. The interaction beam first enters the opposite end of the interaction chamber (not shown). The beam (red) then is focused by a lens through the EOM and is reflected back into the interaction chamber by the retroreflecting mirror (RR). Both the EOM and lens sit on translation stages (TS). For the simpler kicked rotor experiments, we clamped the retroreflecting mirror mount to mounting rods attached directly to the chamber. The EOM and associated optics sit on an irregularly shaped vibration-isolated platform (shaded) next to the chamber. One of the Helmholtz coil mounting forms (a vertical field coil) is shown to illustrate how close to the chamber this platform approaches. The pointed end of the platform is above the corner of the main optical table. One arm of the stimulated Raman beam path is shown (blue) for reference here. The beams are nearly overlapped and some of the optics have had to be cut in half to allow beam clearance. The final quarter-wave plate ($\lambda/4$) in the Raman beam path is in a special mount that allows the other beam to pass very close to the edge of the waveplate.

arrange for phase shifts of over several lattice wells, we have not actually used the EOM to shift the phase more than a single well. In an earlier conception of the state preparation sequence, we imagined ratcheting the phase of the lattice instead of placing the atomic

sample at a nonzero velocity in the laboratory frame. In this scheme, we would sweep the phase uniformly over some period of time and ratchet the phase back to zero at some point when the (amplitude modulated) interaction potential happened to be zero. Ultimately, this scheme turned out to be less than ideal because of the response time of the EOM. A significant problem for common electro-optic materials (for our wavelength) is that they exhibit piezoelectric resonances. These mechanical resonances can be excited by the frequency components of a sudden shift that we apply to the phase. The original EOM that we tested in this location was a Conoptics 350-40 modulator, which has a KD*P (deuterated KDP) crystal. The resonances for this modulator were very severe, and we had significant difficulty changing the phase quickly. The resonances in our lithium tantalate modulator are less severe, but still present. In order to measure the changes in phase, we can temporarily




Figure 4.11: Phase step response of the interaction beam EOM. A small Michelson interferometer was constructed on the EOM platform to measure the phase of the interaction beam as we shifted the control voltage applied to the modulator. Besides the initial transient, there are small oscillations in the phase that are visible for as long as 50 μs after the phase slew begins.

construct a tiny Michelson interferometer on the platform that supports the EOM. A sample phase measurement from the interferometer is shown in Fig. 4.11. The majority of the phase step is completed within $1 \mu\text{s}$, however there is also relaxation on a longer time scale and the phase does not completely stop slewing for about $10 \mu\text{s}$. Besides the slewing, we can see oscillations at the frequency of the piezoelectric resonance, which is near 150 kHz for our modulator. The maximum amplitude of these oscillations is typically of order 1% of the total phase shift. Our SPASM sequence usually involves a period of about $6\text{-}10 \mu\text{s}$ of sliding in the standing wave after we suddenly shift the phase. It is important to note that in this sequence, atomic motion *does* occur while the phase is still changing. Furthermore, there are still minor perturbations to the phase of the standing wave that occur after we have completed our state preparation sequence. We performed simple tests by extending the total state preparation time to allow the oscillations to damp out, however, we were not able to observe any change to the subsequent dynamics. We interpret this to mean that after the first $6\text{-}10 \mu\text{s}$, the phase is approximately constant.

Chapter 5

Chaos-assisted tunneling

5.1 Overview

We have described a state-preparation sequence that allows us to place an atomic distribution at a specific location in phase space. The motivation for these atomic gymnastics is to enable a study of quantum transport in mixed phase space, which we will present in this chapter.

One of the most important quantum mechanical phenomena is tunneling, which we may define loosely as any process through which a particle appears in a region of phase space where it is not classically allowed. (When considering what is “classically allowed,” note that we are neglecting the classical wave-mechanical analogies of our quantum systems.) The most familiar form of tunneling is transmission through an energetic barrier that classical trajectories cannot pass through. Usually we picture the barrier as a classically inaccessible region of space where the potential energy is greater than the total energy of the particle.

An energy barrier is not the only way in which to confine classical trajectories. In some cases, the system dynamics alone can be sufficient to confine trajectories. When examining the phase space structure of the kicked rotor system, it is obvious that the physics of even simple (i.e., simply defined) classical systems can exhibit many types of motion. In some regimes, some or all trajectories are confined to invariant curves (KAM tori). In classical dynamics, a particle on such a trajectory is confined to stay there for all time.

Perhaps a more useful example occurs in the phase space of the simple pendulum (Fig. 1.1). Consider two (noninteracting) particles rotating in this potential with equal

energy $E > E_s$, that is, particles with energy above the separatrix. Let the particles have momenta $p_1(t) > 0$ and $p_2(t) < 0$, so that they are rotating in opposite directions. Suppose that at time $t = 0$, the particle positions are $x_1 = x_2 = 0$. We now have a set of two particles that follow mirror-image trajectories that can be viewed as time-reversed versions of each other. There is no energy barrier in space— in fact, at certain times, the particles occupy the exact same spatial location. And yet, classical transport between these two invariant curves is forbidden. The barrier in this system is implicit in the system dynamics. Tunneling across such a barrier is known as *dynamical tunneling* [Davis81], and is the focus of our study in this chapter. Dynamical tunneling is a general effect, and occurs in many systems where discrete symmetries lead to related but isolated regions of phase space [Chirikov95]. The particular example of dynamical tunneling in a pendulum potential is simply Bragg scattering (§1.4.4).

A slightly more complex situation occurs when we consider dynamics in a system with mixed phase space. We have seen that within certain parameter regimes, the kicked rotor system has large regions of both stable and chaotic motion. In some cases, there are two symmetry-related stable islands that are separated by a chaotic sea. In this case, semiclassical quantization can show that there are quantum states localized upon the stable islands [Tomsovic94].

Let us neglect for a moment the influence of the chaotic region. The system consisting of the two islands is much like the common example of a double-well potential that is used to illustrate tunneling across an energy barrier. The two stable islands constitute wells that can each confine classical trajectories. The quantum dynamics also allow a state to be localized upon the island in the corresponding classical phase space. The symmetry of the system dictates that the eigenstates of the system are not localized in one well or the other, but are symmetric and antisymmetric combinations of the wavefunctions localized in one of the two islands. The antisymmetric state has a slightly higher energy than the symmetric state (by ΔE), with the result that these eigenstates dephase and rephase in time. A quantum state that is localized on one of the two islands is then a linear com-

bination of the two eigenstates. As the eigenstates evolve with respect to each other, the state that began on one island oscillates to the opposite island and back. The tunneling frequency is $\omega_T = \Delta E/\hbar$. This is the essential mechanism for tunneling between these two states, whether tunneling across an energetic or dynamical barrier. There are two potential signatures of this process. First, the (quasi)energy spectrum of the system exhibits a closely-spaced pair (a doublet) of values because the symmetric and antisymmetric states are not degenerate. This is an indirect method of observing tunneling; it is more satisfying if we can directly observe quantum transport by measuring the relevant reversal in position or momentum.

In general, we cannot neglect the influence of chaotic phase space between regions of stability. The complex dynamics in this region lead to possibly irregular quantum states that are localized in this region. These states can overlap with the states that are centered upon the islands of stability and complicate the tunneling process. We can use a semiclassical picture to gain some intuition about the behavior in this case. If the tunneling splitting ΔE is large, it might be easier to tunnel with a multi-step process. We can imagine an atom that is in one of the stable islands can easily tunnel over a small barrier into the chaotic region. Phase space exploration in the chaotic region proceeds rapidly, and visiting the neighborhood of the opposite island may occur rapidly. Once near the other island, the atom can then easily tunnel out of the chaotic region. The behavior of this system can also be described directly in terms of the interaction of a “chaotic” quantum state (one that is localized in the chaotic region) with the two states that occur in the simple tunneling case [Bohigas93; Tomsovic94]. Often there is more than one chaotic state, and in the case that none of them is dominant, the influence of additional states must be considered. Tunneling in this system can occur much more rapidly than would be expected without the chaotic region [Lin90]. This enhancement of the tunneling rate has come to be known as *chaos-assisted tunneling* (CAT) [Bohigas93]. Under small changes in the system parameters, the exact behavior in the chaotic regions may change dramatically. In general, the tunneling splitting (and hence the tunneling rate) tends to fluctuate rapidly as a function of system

parameters [Tomsovic94; Mouchet01].

There have not been many experiments that directly address chaos-assisted tunneling. In the introduction, we noted that a microwave billiards experiment (which is an electromagnetic analogy to a corresponding quantum system) had been used to observe signatures of CAT [Dembowski00]. This experiment observed spectroscopic features (splittings) that show increased tunneling rates in phase space areas that border chaotic regions. Experiments to observe CAT in the context of atom optics have been recently proposed and discussed [Mouchet01; Hug01]. An atom optics experiment similar to ours has recently (i.e., simultaneously with our results) been used to observe dynamical tunneling between stable islands in phase space [Hensinger01]. Note however that dynamical tunneling is not necessarily chaos-assisted, and we must be careful to separate these two concepts. Bragg scattering off of an optical-lattice “grating” is another example of dynamical tunneling that has been observed in atom optics experiments [Martin88; Giltner95; Kozuma99]. Most recently, our experiments have provided clear evidence of chaos-assisted tunneling [Steck01c], which we will describe in this chapter.

We will next introduce the model that we use in our experiment to study chaos-assisted tunneling. The experiment itself requires the state preparation sequence that we have already described, but is otherwise straightforward. The signature of tunneling in our experiment is a reversal of momentum— we begin with atoms at some positive momentum in the lab frame and watch as they flip direction. We are interested in coherent tunneling oscillations between the two islands, and the signature of this behavior is very clear in the experiment. We have performed several tests to verify that tunneling only occurs if certain symmetries are met, as we will see. In addition to these experiments, we have studied the dependence of the tunneling rate on system parameters. Finally, amplitude noise applied to this system can destroy the fragile quantum coherences that allow the tunneling signal.

5.2 Experimental Considerations

5.2.1 Amplitude-modulated pendulum system

The kicked rotor is one of many general models that potentially exhibit chaos-assisted transport. However, it is not a simple system— it is often characterized by interleaved rows of island chains and chaotic regions. It is potentially difficult to identify a set of isolated, symmetry-related islands that lend themselves to the study of CAT. In our experiment, we have studied a slightly different model, the amplitude-modulated pendulum where the phase space is simpler.

These experiments are performed with the same 1D standing wave that we use to form the interaction potential in the kicked-rotor experiments. The chief difference is the form of the amplitude modulation. Rather than turning on the beam in brief pulses, we modulate the intensity of the beam sinusoidally as $I(t) = I_0 \cos^2(\pi t/T)$, where T is the period of the modulation. When we modulate the light, the quantum pendulum Hamiltonian (Eq. 1.46) can be written in physical units as

$$\mathcal{H}(x, p) = \frac{p^2}{2m} - 2V_0 \cos^2(\pi t/T) \cos(2k_L x). \quad (5.1)$$

The quantities here are defined as earlier, and we have shifted the zero of potential energy and the location of $x = 0$. The time-averaged limit of this system is simply the unmodulated pendulum. This is a time-periodic system where the basis states are Floquet states, as in the case of the kicked rotor.

In analogy with §1.4.3 we can simplify the analysis of the system by writing down the Hamiltonian in an appropriate set of scaled units. Let us now define the scaled units for this system as

$$\begin{aligned} x' &= 2k_L x, \\ t' &= t/T, \\ p' &= p\hbar/2\hbar k_L, \\ \alpha &= (\hbar/k)TV_0, \\ \mathcal{H}' &= (\hbar/k)T\mathcal{H}, \end{aligned} \quad (5.2)$$

where the scaled Planck constant is again identified as $\hbar = 8\omega_r T$. We now rewrite Eq. 5.1 in terms of the scaled units, and drop the primes to find the dimensionless Hamiltonian

$$\mathcal{H}(x, p, t) = \frac{p^2}{2} - 2\alpha \cos^2(\pi t) \cos x. \quad (5.3)$$

The scaled well depth α acts much like the stochasticity parameter in the kicked-rotor system and specifies the classical dynamics completely.

Surfaces of section for this system are shown with several values of α in Fig. 5.1. This system has three primary resonances, which we may see by rewriting the Hamiltonian as

$$\mathcal{H} = \frac{p^2}{2} - \alpha \cos x - \frac{\alpha}{2}(\cos(x + 2\pi t) + \cos(x - 2\pi t)). \quad (5.4)$$

The potential is then equivalent to the sum of three superimposed pendulum potentials, one at zero velocity, and two moving in opposite directions with velocity $\pm 2\pi$. For low (but nonzero) values of α , the system has three clear resonances that are centered at $x = 0, p = 0, \pm 2\pi$. As we turn up the strength of the potential, the regions where the resonances overlap break down into layers of chaos. The stable islands centered at the locations of the resonances survive over a wide range of α .

The two islands that are centered at nonzero momentum comprise the symmetry-related pair that we use to study tunneling. In the experiment, we place atoms in one of the two islands and observe tunneling to the opposite one. These islands are (for some values of α) connected by a wide chaotic region, and so we expect chaos-assisted tunneling to play a role in certain parameter regimes. The island at $p = 0$ is not coupled by symmetry to the outer islands, and does not play a substantial role in the tunneling process. For the values of α where most of our experiments are conducted (near $\alpha = 10$) the center island is barely visible.

5.2.2 Symmetry requirements

As we have stressed, tunneling is facilitated by symmetry. The most important symmetry is with respect to reflection across $p = 0$. This is equivalent to a time-reversal symmetry of




Figure 5.1: Surfaces of section for the driven pendulum with various α . The dynamics of this system are given by Eq. 5.3. The values of α are 0.4 (upper left), 4.0 (upper right), 10.4 (lower left) and 13.6 (lower right). The system exhibits three primary resonances, one at $p = 0$, and two centered at opposite momentum values. As the value of α is increased, the system becomes more chaotic. The resonance at $p = 0$ is not visible at the higher values of α shown here, but does become visible again if it is increased further. Most of our experiments were performed near $\alpha = 10$, where two clear islands are visible. These surfaces of section show the phase space as sampled every two modulation periods. (Surfaces of section by D. Steck.)

the system [Casati94]. Even minor deviations from this symmetry are expected to inhibit tunneling [Chirikov95; Tomsovic98]. A basic requirement to observe CAT is that the atoms are actually loaded into one of the relevant islands. Besides giving us a well-defined initial

condition, this ensures that the mirror-image section of phase space has the same regular structure. A secondary requirement for symmetry is imposed by the quantum nature of the system. As we seen in §1.4.1, atoms in our experiment only exchange momentum with the interaction beam in units of $2\hbar k_L$ (\tilde{k} in scaled units). A consequence of this is that atoms are only coupled to their symmetric reflections across $p = 0$ if their momentum is an integer multiple of $\hbar k_L$ ($\tilde{k}/2$). Suppose that an atom has momentum $p = 4\tilde{k}$. The mirror-image momentum state has momentum $p = -4\tilde{k}$, and these two states are coupled by an integer number (8) of two-photon hops. If the initial momentum were instead $p = 4\tilde{k} + \epsilon$, then the symmetric reflection is at $p = -4\tilde{k} - \epsilon$. The momentum difference between these two states is $8\tilde{k} + 2\epsilon$, so the two states are not coupled by the standing wave unless ϵ is an integer multiple of $\tilde{k}/2$. In order to have true symmetry in this system, we must therefore meet the very stringent requirement of populating only “integer” momentum states. This is the same symmetry requirement that must be met for Bragg scattering (§1.4.4). In order to observe chaos-assisted tunneling, we must also load the atoms into the proper island in phase space. This means that we must be careful to pick a parameter regime where the classical island includes a momentum value that is on the proper momentum ladder. These requirements of placing atoms with arbitrarily narrow momentum at a very specific location in phase space necessitate the use of our entire state preparation sequence.

5.2.3 CAT Experiments

The experimental setup and sequence for these experiments is that which we have described in great detail while describing the state preparation sequence (§4.4). Most of our experiments have been conducted with a well depth chosen such that α is near 10. As we have seen in Fig. 5.1, this is a value where the phase space has two clear islands. For most of the experiments here, the modulation period T was chosen to be $20 \mu\text{s}$, which yields a scaled Planck constant $\tilde{k} = 2.08$. For a few experiments, we used a $10 \mu\text{s}$ modulation period so that $\tilde{k} = 1.04$. We used our state preparation sequence with an $800 \mu\text{s}$ Raman tagging pulse to prepare a compact initial condition and center it upon the location of the island. The momentum ladder defined by the Raman pulse must be carefully chosen so that the

atoms have the proper reflection symmetry. In practice, we pick the momentum ladder by loading atoms into the island and scanning the tagging frequency (which defines the momentum ladder) over a small range until we see atoms tunneling. An illustration of our initial conditions is shown in Fig. 5.2.

Let us review some of the parameters that were used in these experiments and in their state-preparation stages. For the experiments with $\tilde{k} = 2.08$, the island location for $\alpha = 10.5$ was slightly above $4 \cdot 2\hbar k_L$. In order to place atoms near the island, we used the SPASM sequence with a $6 \mu s$ “sliding” period to accelerate atoms. The well depth for the




Figure 5.2: The phase space for a typical CAT experiment, as given by Eq. 5.3 with $\alpha = 10.5$. Also shown is an illustration of our initial conditions (for $\tilde{k} = 2.08$), visible here as narrow red stripes near the upper island. The “striped” structure of the initial condition is much like the schematic representation in Fig. 4.7. (Figure: D. Steck)

sliding period was picked so that the final momentum distribution was peaked at $4.1 \cdot 2\hbar k_L$, with a width of $\sigma_p/2\hbar k_L = 1.1$. In a second data set at $\tilde{k} = 2.08$, we swept the value of α over a wide range and introduced amplitude noise into the system. For these cases, we used a deeper well depth for the state preparation and accelerated the distribution to $4.2 \cdot 2\hbar k_L$ with a $4.5 \mu\text{s}$ sliding period. The final width of the distribution envelope was $\sigma_p/2\hbar k_L = 1.7$ in this case. For the experiments with $\tilde{k} = 1.04$, we used a very deep lattice to accelerate the atoms to $8.2 \cdot 2\hbar k_L$, with a width $\sigma_p/2\hbar k_L = 2.1$. For this value of \tilde{k} , the islands appear near $8 \cdot 2\hbar k_L$.

For most of the data that we present, the atomic distribution was sampled every two modulation periods, which corresponds to the sampling for the phase space portraits in Fig. 5.1. For a few of the cases (shown in §5.3.5), we sampled at a much faster rate to observe the dynamics more closely.

In each of the experiments we measured the momentum distribution after some interaction time. In many cases, we see a substantial number of atoms tunnel to the momentum opposite that which they began with. In the data shown here, we usually reduce the momentum distribution to the single parameter $\langle p \rangle$, the average momentum value. Although our initial momentum for the experiments is often centered near (say) $4.1 \cdot 2\hbar k_L$, the value of $\langle p \rangle$ may appear substantially lower. Much as in the determination of the absolute ensemble energy of an atomic distribution (See §3.6.2), it is not practical to compensate for most of the systematic factors that affect this value. It is still a good measure of the degree of tunneling, however, so long as we pay attention to *oscillations* in the value of $\langle p \rangle$, as only tunneling processes will affect $\langle p \rangle$ in an oscillatory manner.

Finally, we have noted that in the state preparation sequence, we end up tagging a very small number of the initial atoms. Nearly all of the data presented in this chapter have been averaged over many individual experiments to improve the signal to noise ratio.




Figure 5.3: Observation of chaos-assisted tunneling oscillations. Momentum distributions are measured at $40 \mu\text{s}$ intervals (twice the modulation period) for atoms that are initially loaded into the island near $p = 4 \cdot 2\hbar k_L$. A substantial fraction of the atoms coherently tunnel to and from the island at opposite momentum. This data corresponds to the initial conditions and phase space shown in Fig. 5.2. Each slice of this data represents 20 measurements that were averaged together to improve the signal-to-noise ratio.




Figure 5.4: Individual momentum distributions corresponding to the first four highlighted distributions from Fig. 5.3. These distributions are averaged over 100 iterations of the experiment to increase the signal-to-noise ratio.




Figure 5.5: Observation of CAT oscillations with $k = 1.04$ and $\alpha = 11.2$. The modulation period for this data set is $10 \mu s$ (to realize $k = 1.04$), so the data are sampled every $20 \mu s$. These data are averaged over 10 iterations of the experiment.

5.3 Data and Results

5.3.1 Observation of tunneling

The first chaos-assisted tunneling data from our experiments are shown in Fig. 5.3. This data was taken with $\alpha = 10 \pm 5\%$, $k = 2.08$, and corresponds to the initial conditions and phase space of Fig. 5.2. In this data, about half of the initial distribution is observed to tunnel to the opposite momentum. Four oscillations are observed in this data. Initially, we believed that the damping could be an effect of decoherence in the system. Some of the later data that we have taken shows that we can observe tunneling for a much longer time than is visible here— it is most likely that the loss in contrast here is the result of beating between two tunneling frequencies. Several individual momentum distributions that correspond to this data are shown in Fig. 5.4. Again, about half of the atoms apparently tunnel across $p = 0$. We believe that the reason that more atoms do not tunnel is the residual width of our momentum distribution.

We have not at this point proven that these oscillations are in fact chaos-assisted

tunneling oscillations. Even the simple reversal of momentum that we observe could potentially be explained by other mechanisms, especially considering that we are not observing *all* of the motion, but rather a stroboscopically sampled portion of it. Classical dynamics, such as motion in the pendulum inside the separatrix, can periodically reverse the momentum of an atomic sample. If some of the initial condition were just outside the separatrix, it would remain at positive momentum. While this is a plausible scenario, there are many signs that this is not the case in our experiment. The amplitude-modulated potential that our experiment employs is well defined, as is the initial condition. Shortly, we will verify the structure of the classical phase space by showing that tunneling only occurs when the atoms are placed directly on the island. We can demonstrate the quantum nature of the system and prove that this is in fact dynamical tunneling, by showing that the system is highly sensitive to the exact momentum state that is chosen. Showing that we have chaos-assisted tunneling is not as simple. There are several key observations, each of which is consistent only with chaos-assisted processes. Most importantly, the tunneling that we observe occurs on a timescale very fast compared to that which would be expected in a fully integrable system. There is also supporting evidence in the tunneling rate as a function of system parameters, and in the details of the motion from one island to the other.

An example of chaos-assisted tunneling in a different regime is shown in Fig. 5.5. Here, we observe tunneling with $\tilde{k} = 1.04$ and $\alpha = 11.2$. As we have stated, the initial condition for this data is near $p = 8 \cdot 2\hbar k_L$, and so the tunneling is a 32-photon transition. A similar fraction of the atoms tunnel in this case, which is presumably limited by the residual width of the momentum distribution. In practice, it is quite difficult for us to change the value of \tilde{k} because the new momentum scale requires that we recalibrate the full SPASM sequence.

5.3.2 Sensitivity to island overlap

When we described the state preparation sequence, we mentioned that we can translate our initial wavepacket to locations other than $x = 0$ simply by allowing a free drift period before




Figure 5.6: Initial conditions at other locations in phase space. Schematic representation of our initial conditions in phase space for experiments in which we displace the initial conditions spatially by allowing free drift before the experiment begins. The drift times that are shown here correspond to the data shown in Fig. 5.7. An ellipse drawn around the outline of the distribution is meant to guide the eye. (Figure: D. Steck)

we turn on the interaction potential. This works because the initial condition that we have generated has a nonzero net velocity. The results of the free-drift procedure are illustrated in Fig. 5.6, where we see the initial condition at other locations in phase space. In the cases shown here, we allowed the distribution to drift across $1/4$ of a lattice period, $1/2$ of a lattice period, and over a full lattice period. The time for the peak of the distribution to drift across an entire lattice period is about $15 \mu\text{s}$.




Figure 5.7: Sensitivity of tunneling signal to the initial spatial location of the atomic distribution. The average momentum $\langle p \rangle$ is measured as a function of time for $\alpha = 10.5$ and $\tilde{k} = 2.08$ for several displacements in x . These displacements are controlled by the drift time after the state preparation sequence and are illustrated in Fig. 5.6. The four cases shown here are without drift (filled circles), $3.8 \mu\text{s}$ drift time (squares), $7.6 \mu\text{s}$ drift time (triangles), and $15.1 \mu\text{s}$ drift time (open circles). The case without drift corresponds to the data shown in Fig. 5.3. These data represent averages over 20 individual experiments.

The data for these different drift times is shown in Fig. 5.7. If we displace the initial condition by $1/4$ of a lattice period, the tunneling is significantly inhibited, but still visible. The slices of the wavepacket that we have drawn show the 50% contours of the atomic distribution— in the $1/4$ period drift case, there is still significant overlap with the island. If we allow the atoms to drift across $1/2$ of the lattice period, they are as far away from the stable island as they can get. The average momentum $\langle p \rangle$ in this case does not reveal any tunneling. The initial condition appears to be in a regime where the motion is regular, much like the motion above the separatrix in the unmodulated pendulum potential. It is likely that tunneling in this case *could* be observed, but it probably would not be a form of chaos-assisted tunneling. For a $15.1 \mu\text{s}$ drift time, the atoms drift across a full lattice period and are again centered upon the island of stability. In this case we again observe tunneling oscillations, although their magnitude is decreased, presumably because of the distortion of

the wave packet.

The purpose of these tests is partially to verify the classical dynamics of the system. If there were some other transport mechanism that would bring atoms to $p < 0$, we would expect to see it whether or not the initial distribution were centered on the island. Beyond this simplistic test, the chaos-assisted tunneling mechanism that we have described requires projecting our initial distribution into a state that is localized upon the island of stability. The fact that tunneling is suppressed when we are not initially on the stable island confirms that the tunneling process that we observe is between the islands.

5.3.3 Sensitivity to velocity class

As we have discussed, the requirements for true symmetry in our system are very stringent. In order for atoms to be coupled to their symmetric reflections across $p = 0$, they must essentially comprise a plane wave with a specific, half-integer momentum. Our initial conditions for experiments with $\tilde{k} = 2.08$ are peaked near $p = 4 \cdot 2\hbar k_L$, so that the most




Figure 5.8: Sensitivity of CAT to changes in momentum ladder, for $\alpha = 10.5$ and $\tilde{k} = 2.08$. In these cases, we selected velocity groups that were centered at $p = 0$ (circles), $p = 0.05 \cdot 2\hbar k_L$ (squares), and $p = 0.12 \cdot 2\hbar k_L$ (triangles). Even modest changes in this initial velocity class (compared to $\hbar k_L$) are sufficient to suppress tunneling.

populated integer momentum state is at $4 \cdot 2\hbar k_L$. In the experiment we have been able to observe tunneling if the initial momentum is instead 3 or 3.5 double recoils. This is a coarse statement about where on the momentum ladder we are operating, rather than a statement about how the momentum ladder is defined. In this sense, coarse (half-integer) displacements in momentum have the effect of moving the initial condition off of the island. As with the displacements in position, we expect that the tunneling signal will be attenuated if the initial condition does not overlap with the island.

Rather than the coarse position of our distribution, we are now concerned with the exact position and residual width of the momentum slices within the distribution. The momentum class that is tagged depends upon the detuning between the relevant frequency




Figure 5.9: Sensitivity of CAT to width of momentum slices in the distribution, for $\alpha = 10.5$ and $\tilde{k} = 2.08$. Our most common momentum slice was derived from an $800 \mu s$ stimulated Raman tagging pulse (circles), which produces a momentum width (HWHM) of $0.03 \cdot 2\hbar k_L$. A slightly wider initial momentum width is generated by a $400 \mu s$ pulse, yielding a HWHM of $0.06 \cdot 2\hbar k_L$ (squares). A $200 \mu s$ pulse selects an initial width (HWHM) of $0.12 \cdot 2\hbar k_L$ (triangles). Also shown is a case where we did not apply any tagging pulse at all, but merely employed the 3D lattice cooling (heavy solid line). In this case, the distribution has a HWHM $0.8 \cdot 2\hbar k_L$. These data are averaged over 20 ($800 \mu s$ tag), 10 ($400 \mu s$ tag), and 5 ($200 \mu s$ tag) iterations for the tagged cases. The lattice-cooled curve is the data for a single sweep of the experiment and is not averaged.

components of the stimulated Raman beams. Usually, we select this detuning such that the tagged atoms are at $p = 0$, so that we may see a large tunneling signal. This initial momentum is the value immediately after tagging, which defines the momentum ladder for the initial condition for the experiment. If we vary this detuning, we can select a momentum ladder that is no longer centered at $p = 0$. This effectively breaks the symmetry of the system, and we expect it to inhibit tunneling.

We have experimentally studied the effects of choosing different momentum ladders, and the data are shown in Fig. 5.8. A shift in the location of the momentum ladder by $0.12 \cdot 2\hbar k_L$ is enough to nearly suppress the tunneling altogether. This is a very small displacement when compared with the $2\hbar k_L$ spacing of the stripes in the momentum distribution, or as compared to the average initial momentum. For a displacement of $0.05 \cdot 2\hbar k_L$, the tunneling




Figure 5.10: CAT experiment without Raman velocity selection. We repeated the experiment that was shown in Fig. 5.3, except that we used lattice cooling alone without a Raman velocity selection stage. Only a small fraction of the distribution appears near the opposite island, and coherent oscillations (the signature of tunneling in this system) are not observed.

signal is only slightly attenuated. Not coincidentally, this displacement is similar to the initial width of the distribution. This data was taken with $800 \mu\text{s}$ velocity selection pulses so that the tagged atoms have a momentum width (HWHM) of $0.03 \cdot 2\hbar k_L$.

Our next task is to examine the sensitivity to the remaining momentum width of our momentum slices. It is not practical in our experiments to perform tagging pulses that are much longer than the $800 \mu\text{s}$ pulses that we use. This places a limit upon how narrow our momentum slices can be. While we cannot further decrease the width of our momentum slices, we can easily *increase* the width. Shorter Raman tagging pulses result in wider momentum slices. The result of this is that proportionately fewer of the atoms in the distribution are on the correct momentum ladder to exhibit tunneling. The contrast of the tunneling signal is then expected to decrease as a wider velocity class is selected. Data from our experiment showing this dependence appear in Fig. 5.9. Also shown in this figure is the signal that we get if we simply use 3D lattice cooling, without a stimulated Raman tag. The initial distribution in this case is still localized in the correct region of phase space, but the momentum distribution is no longer composed of stripes, and the outline of the wavepacket is also enlarged. In this case, the atoms that tunnel represent a very small fraction of the overall sample, as we can tell from the momentum distributions shown in Fig. 5.10. This is sufficient to show that chaos-assisted tunneling in this system cannot be observed without a subrecoil initial distribution. The residual width of our initial distribution turns out to be enough to explain the incomplete tunneling that we see in our experiment. More complete tunneling could be achieved with a colder yet initial condition such as an atomic Bose-Einstein condensate (BEC). A sodium BEC was used to observe more complete tunneling in another group's dynamical tunneling experiments [Hensinger01].

5.3.4 Comparison with pendulum

We have asserted that tunneling in our system occurs more rapidly than would be expected from a simple two-state tunneling process. If we simply time-average our potential we instead realize a simple pendulum, which is our prototypical integrable system. Our initial

conditions in the pendulum phase space are illustrated in Fig. 5.11, for an average well depth corresponding to that of modulated experiments with $\alpha = 10.5$ and $k = 2.08$. Our initial condition is (almost entirely) above the separatrix and therefore constrained to remain above $p = 0$ by the classical dynamics. The dominant transport mechanism to $p < 0$ in this system is then high order Bragg scattering. The symmetry requirements for Bragg scattering are the same as those for chaos-assisted tunneling, so we know that we have met the general requirements for Bragg scattering to occur. Bragg scattering is a form of




Figure 5.11: Initial conditions in the corresponding pendulum phase space. This figure depicts the pendulum phase space that is realized by averaging (over time) the amplitude-modulated pendulum potential that is represented in Fig. 5.2. In the experiment that this represents, we begin with the same initial conditions that we use to observe CAT, which are shown here. The momentum axis for this particular phase space is calibrated in the scaled units of the quantum pendulum.

dynamical tunneling that can occur in the absence of any chaotic region of phase space, and provides a basis for studying the effects of the chaotic region of phase space.

When we place atoms in the corresponding static pendulum potential, we do not observe any tunneling. The data from the experiment are shown in Fig. 5.12, where motion in the modulated and unmodulated pendulum potentials are compared. Our initial conditions are near $4 \cdot 2\hbar k_L$, so the relevant tunneling process in the static pendulum case is 8th order Bragg scattering. For our parameters, the expected Bragg period is roughly 1 s, and so we do expect to see this process over the timescale of our experiments. The typical period of our CAT oscillations in this regime (where $\bar{k} = 2.08$) is $400 \mu\text{s}$, several orders of magnitude faster than the Bragg process. We have also seen in Fig. 5.5 that we can observe oscillations when $\bar{k} = 1.04$. In this case, our initial momentum state is near $p = 8 \cdot 2\hbar k_L$, so the tunneling is occurring across the 32-photon gap of $p = \pm 8 \cdot 2\hbar k_L$. This tunneling occurs in our experiments with a period of about $200 \mu\text{s}$. We did not explicitly verify that the unmodulated system does not display tunneling, but the corresponding 16th order Bragg




Figure 5.12: Comparison of chaos-assisted tunneling oscillations (circles) to motion in the corresponding quantum pendulum potential (squares). The CAT experiment was conducted with $\alpha = 10.5$ and $\bar{k} = 2.08$. In the unmodulated pendulum case, we used the time-averaged intensity of this potential. These data are averaged over 20 iterations per point.

scattering is expected to have a period of about 20 years.

It is slightly unfair to compare our results directly with Bragg scattering from an unmodulated pendulum potential, because we are neglecting the possibility that the modulation enhances the probability of *direct* tunneling between the two islands of stability. This is as opposed to tunneling enhanced by interaction with additional (chaotic) states. The Bragg system is much like a driven two-level atomic system. In such a system, the time dependence of the driving field can enhance the transition probability [Allen87]. The enhancement factor for the $\tilde{k} = 2.08$ case is about 50, and for $\tilde{k} = 1.04$ it is about 9000 [Steck01b]. These enhancement factors are still far below what would be required to explain our experimental results, and thus provide supporting evidence for chaos-assisted tunneling.

5.3.5 High temporal resolution data

The data that we have presented thus far have been sampled once after every two modulation periods. This sampling is convenient for many reasons, most particularly because the islands in phase space are stationary when sampled at this period. Furthermore, stroboscopic sampling allows us to save a great deal of time in the experiment when we observe long-time dynamics. The continuous dynamics of the system are also interesting, and can be studied by sampling the data at much shorter intervals. As we will see, the dynamics on shorter time scales can reveal behavior that further indicates the presence of chaos-assisted tunneling mechanisms.

When we plot (stroboscopically) the phase space of our amplitude modulated pendulum at phases of the modulation other than zero, we can see what happens to the islands over the course of a single modulation period. Phase space portraits sampled at other phases of the modulation are shown in Fig. 5.13, for our typical case of $\alpha = 10.5$. These phase space portraits are technically Poincaré surfaces of section, which are two-dimensional slices of a three-dimensional phase space. In the sections through the phase space at other times, we see that the classical islands rotate through the phase space, moving out in momentum as they avoid the remnants of the (no longer visible) central island. We have previously




Figure 5.13: Classical space portraits at different phases of the modulation. These surfaces of section are sampled stroboscopically with the phase offsets indicated here for the amplitude-modulated pendulum with $\alpha = 10.5$. The islands move towards the momentum axis, but do not cross it. The continuous behavior of classical particles trapped in the upper island is shown in Fig. 5.14.

considered only the sampling at zero phase offset, where the islands are at their maximum distance from the momentum axis. When we sample the phase halfway through the modulation period, the islands have moved closer to the momentum axis (which they do *not* cross).

As a minor diversion, note that the surface of section with a sampling phase offset of $T/2$ is the same as if we had modulated the system as $\sin^2(t)$, without a detection phase offset. Along the lines of the tests in §5.3.2, we verified experimentally that we can observe




Figure 5.14: Motion of the classical island in momentum space. In this classical simulation, an ensemble of particles was loaded into the island of stability that has positive momentum. The motion of these particles is traced as the sampling phase is changed. The momentum of the particles trapped in the island changes to reflect the island location. Note that these trajectories vary in momentum, but are confined to positive momentum. For this simulation, $\alpha = 10.5$, and the motion of the island corresponds to the phase space portraits shown in Fig. 5.13.

CAT when the amplitude is modulated as $\sin^2(t)$. In this case, the islands appear centered at $x = \pm\pi$, at a slightly lower momentum than for the $\cos^2(t)$ case. We loaded the atoms into this island by beginning at a lower initial momentum (near $p = 3 \cdot 2\hbar k_L$), and allowed a free drift period so that they drifted halfway across the phase space to be centered on the islands. The tunneling oscillations for this system were nearly as clean as in our standard $\cos^2(t)$ experiments.

When we load atoms into an island in this system, it is important to realize that the position of this island does vary in time. This is not an arbitrary artifact of the sampling time, but rather is an effect of the system dynamics. In order to see this directly, we can




Figure 5.15: High temporal resolution tunneling, $k = 2.08$, $\alpha = 7.7$. This data set shows the momentum distribution in our experiment over the time scale of the first chaos-assisted tunneling oscillation. The momentum distributions are sampled every $2 \mu\text{s}$, which is ten times per modulation period. The classical motion of the islands is visible as an oscillation that repeats after each modulation period. These distributions represent averages over 10 individual experiments.

monitor the phase-space location of particles in the island as the sampling phase is changed. In Fig. 5.14 we show the momentum of particles trapped in the island as the sampling phase is changed, corresponding to momentum changes on timescales below that of the modulation.

It is also possible to directly observe these oscillations in our experiment. In Fig. 5.15, Fig. 5.16, and Fig. 5.17, we present measurements of the atomic momentum distribution where the distribution is sampled at a rate of 10 times per modulation period. These data are collected over a total time representing one full tunneling period, rather than over many tunneling oscillations. The data in these figures show motion at several time scales. First, there is the timescale of the tunneling, of which we can see one complete oscillation during each scan which is 20 modulation periods long. The motion of the classical




Figure 5.16: High temporal resolution tunneling, $k = 2.08$, $\alpha = 11.2$. This data set shows the momentum distribution in our experiment over the time scale of the first chaos-assisted tunneling oscillation. The momentum distributions are sampled every $2 \mu\text{s}$, which is ten times per modulation period. The classical motion of the islands is visible as an oscillation that repeats after each modulation period. These distributions represent averages over 10 individual experiments.

islands is cyclic with the same period of the modulation. This motion is most clearly visible in the first couple of modulation periods, before a significant amount of tunneling occurs. After some tunneling has occurred and an appreciable population is visible in the island at negative momentum, it is possible to see these atoms oscillate in the opposite islands.

A secondary feature that we can see in this data is the momentary appearance of atomic population in the wide region between the two islands. Atoms appear to tunnel to the opposite island gradually, passing through the chaotic region but not building up there. This sort of behavior is an independent signature of chaos-assisted tunneling, as it directly suggests that the tunneling is made possible by the interaction with states localized in the chaotic region [Tomsovic94].




Figure 5.17: High temporal resolution tunneling, $k = 1.04$, $\alpha = 10.5$. This data set shows the momentum distribution in our experiment over the time scale of the first chaos-assisted tunneling oscillation. The momentum distributions are sampled every $1 \mu\text{s}$, which is ten times per modulation period. The classical motion of the islands is visible as an oscillation that repeats after each modulation period. These distributions represent averages over 10 individual experiments. Note that the signal to noise ratio for this data is much lower than for the data at $k = 2.08$, because we are imaging a much larger spatial region to measure the wider momentum distribution. The coarse horizontal stripes visible on this data are an artifact of the measurement system, usually hidden by larger signal levels.

5.3.6 Sensitivity to well depth

We have already seen several signatures in our experiment that indicate the presence of chaos-assisted tunneling. One additional signature that we expect is a sensitivity to the parameters of the system. The parameter that we can easily sweep is the value of α , which is controlled by the laser intensity during the interaction phase of the experiment. (We can also change the value of k , but it is not practical to sweep this parameter.) In direct (two-state) tunneling between the islands, we would expect the tunneling rate to increase as a function of α . However, in chaos-assisted tunneling there is no universal parameter dependence.




Figure 5.18: CAT sensitivity to α , $\hbar = 2.08$. The value of $\langle p \rangle$ is plotted as a function of time for a range of α values that includes the region where we see strong tunneling oscillations. The average momentum $\langle p \rangle$ is indicated by color, where lighter colors correspond to higher values of $\langle p \rangle$. Each point in this data set is an average over 5 individual experiments.

We have measured the average momentum $\langle p \rangle$ as a function of time for a wide range of α for $\hbar = 2.08$, as shown in Fig. 5.18. A tunneling signal is visible (as an oscillation in $\langle p \rangle$) only within the range $\alpha \approx (7 - 14)$. The classical phase space portraits in Fig. 5.1 can explain part of this range. At the upper end of this range, the outer islands in phase space destabilize and tunneling is inhibited. At the lower end of the range, there is not such an obvious change in the classical dynamics. Although the central island does become larger




Figure 5.19: CAT oscillations with $\alpha = 8.0$, $k = 2.08$. This data is one slice from Fig. 5.18. In this data set we can see coherent oscillation for the entire timescale of the experiment.




Figure 5.20: CAT oscillations with $\alpha = 9.7$, $k = 2.08$. This data is one slice from Fig. 5.18. At this value of α , there are two tunneling frequencies that beat against each other, producing the structure that we see here.

as α is decreased, it does not divide the phase space until α is much lower. A possibility is that the tunneling rate abruptly becomes much lower than we are able to observe in the




Figure 5.21: Tunneling rate as a function of α , $k = 2.08$. These tunneling periods were numerically extracted from the data shown in Fig. 5.18. For α in the range 8.9 to 10.3, the tunneling rate is fitted to two independent frequencies. The points of zero frequency are points where the fitting procedure fails when there is not a clear oscillation frequency.

experiment. It is apparent that at least one of the conditions that enable chaos-assisted tunneling is abruptly broken at some value of α .

There are several other features to note in the data shown in Fig. 5.18. First, notice that at certain values of α , the oscillations remain coherent for at least the 3.6 ms length of the experiment. The data for of these cases ($\alpha = 8.0$) is shown for clarity in Fig. 5.19. At some other values of α , there is more than one frequency involved in the tunneling process. We can see this directly in the $\langle p \rangle$ curve for $\alpha = 9.7$, presented in Fig. 5.20. In this data, we can see two tunneling frequencies beating against each other. The tunneling rates for the entire data set are shown in Fig. 5.21, where it is apparent that the second tunneling frequency is only visible in a much smaller range of α . The main oscillation frequency that is visible in this plot exhibits a decreasing trend as α is increased, which is clear evidence that the tunneling mechanism is not direct tunneling. Within the range of tunneling, we do not see any rapid changes in the tunneling rate as a function of α . If there were very rapid changes, it is likely that we would not see them because of systematic effects that cause

minor averaging over the value of α . These effects are analogous to those that we discussed in the context of the kicked rotor that lead to averaging over the stochasticity parameter K .

For $\hbar = 1.04$, tunneling occurs over a more narrow range of α , as we see in Fig. 5.22. The tunneling rates for this data set are presented in Fig. 5.23. From these plots, there is no significant overall trend in the tunneling rate. There may be some significant changes in




Figure 5.22: CAT sensitivity to α , $\hbar = 1.04$. The value of $\langle p \rangle$ is plotted as a function of time for a range of α values that includes the region where we see strong tunneling oscillations. The average momentum $\langle p \rangle$ is indicated by color, where lighter colors correspond to higher values of $\langle p \rangle$. Each point in this data set is an average over 10 individual experiments.

the tunneling rate at the edges of the tunneling range, but our experiment does not have enough resolution to resolve these fluctuations. It is worth noting that the step size in the sweeps of α was chosen to be close to that of the stability of α over a data run, which is about $\pm 1\%$. A more detailed study of these dynamics might then be facilitated by a more stable laser system.




Figure 5.23: Tunneling rate as a function of α , $k = 1.04$. These tunneling periods were numerically extracted from the data shown in Fig. 5.18. The points of zero frequency are points where the fitting procedure fails when there is not a clear oscillation frequency.

5.3.7 Sensitivity to amplitude noise

Earlier, we studied the effects of amplitude noise on the kicked rotor system and found that for sufficiently strong noise, we could drive the system to the classical limit. The kicked rotor is a particularly interesting system in which to study correspondence because the classical and quantum limits are characterized by easily distinguishable behavior. In that system, the amplitude noise destroys the fragile quantum coherences that cause dynamical localization. Chaos-assisted tunneling is also an effect of quantum coherence and therefore we expect it to also be susceptible to the effects of noise [Grossmann93].




Figure 5.24: Modulated pendulum intensity profile with and without amplitude noise. We measured the intensity of the interaction beam for the latter part of the SPASM sequence and the first part of the interaction time, where the optical lattice is modulated as $\cos^2(t)$. In the first part of these curves, we see the interaction beam intensity steadily increasing. This is the end portion of the adiabatic deepening of the lattice that we use to load atoms into the standing wave wells. Once the lattice is fully bright, we hold the intensity constant and shift the phase so that the atoms accelerate in the lattice. After this acceleration phase, we begin the interaction phase of the experiment. The period of the modulation is $10 \mu\text{s}$, which is what we use to obtain $\hbar k = 1.04$ in the experiment. The two cases shown here represent zero added amplitude noise (dashed) and 15.7% RMS deviation amplitude noise (solid). The instantaneous deviation of the noise is proportional to the instantaneous amplitude of the non-noisy case.

We added amplitude noise to our modulated pendulum system by replacing the Hamiltonian Eq. 5.3 with

$$\mathcal{H} = \frac{p^2}{2} - 2\alpha[1 + \gamma(t)] \cos^2(\pi t) \cos(x) . \quad (5.5)$$

The instantaneous noise factor $\gamma(t)$ randomly fluctuates about zero with a symmetric distribution. This distribution is generated digitally on the computer that programs the arbitrary waveform generator that controls the interaction beam. We begin with a Gaussian white noise sample at 10 MHz. This signal is then sent through a digital Chebyshev low-pass filter to remove high frequency components on the noise. For the $k = 2.08$ case, the filter cutoff frequency was 500 kHz, and for $k = 1.04$ the cutoff frequency was 1 MHz. The filtering removed most of the large deviations in the noise signal. In some cases with strong noise, the instantaneous intensity of the light had to be clipped at the output range of the laser (most of the clipping occurs at zero intensity—our laser cannot go negative). Clipping was only significant for our strongest noise case, “62%” noise. For this system, we specify the




Figure 5.25: Effects of applied amplitude noise on the tunneling oscillations for $\alpha = 11.2$ and $k = 2.08$. The RMS noise levels are 0% (circles), 15.7% (squares), 31% (diamonds), and 63% (triangles). The tunneling is only completely suppressed at the 63% level, and thus is substantially less sensitive than in the $k = 1.04$ case in Fig. 5.26. The data are averaged over 10 realizations of noise, and are sampled every 2 modulation periods.

level of amplitude noise by the RMS noise level $\langle \gamma^2(t) \rangle^{1/2}$, which is determined *after* the filtering procedure. An example of the intensity profile corresponding to Eq. 5.5 is shown for one particular realization of noise in Fig. 5.24.

The original purpose of the filtering procedure was to facilitate comparison with measurements of phase noise effects. Unfortunately, the piezoelectric resonances in our electro-optic phase modulator have made it impossible to study even such a bandwidth-limited realization of phase noise. Nevertheless, the low pass filter is useful because it keeps the amplitude noise changes within the bandwidth of the AOM that modulates the interaction beam. If we did not low-pass the noise signal in a controlled way, it would be smoothed by the electronic and optical systems in ways that might significantly affect the effective noise level.

The effects of amplitude noise on our tunneling signals for $k = 2.08$ are shown in Fig. 5.25. For this data set, $\alpha = 11.2$. When we do not apply any amplitude noise to




Figure 5.26: Effects of applied amplitude noise on the tunneling oscillations for $\alpha = 11.2$ and $k = 1.04$. The RMS noise levels are 0% (circles), 7.9% (squares), 15.7% (diamonds), and 31% (triangles). The tunneling is completely suppressed at the 31% level, and thus is more sensitive than in the $k = 2.08$ case in Fig. 5.25. The data are averaged over 10 realizations of noise, and are sampled every 2 modulation periods.

the system, we see oscillations in the value of $\langle p \rangle$ as we have seen before. For even small amounts of added noise, the value of $\langle p \rangle$ begins to relax towards zero at longer times. This effect is partly due to the way that the noise affects the classical dynamics of the modulated pendulum system. We expect the noise to cause diffusion in momentum that can bump particles out of the classical islands. After this happens, tunneling is not expected to be significant as the particles diffuse throughout the chaotic region of phase space. As the noise level is increased, the range of time over which we can observe tunneling decreases. For the noisiest case in this data set, we do not observe even one tunneling oscillation.

We have also studied the effects of noise at $\hbar = 1.04$, shown in 5.26. This data was also taken with $\alpha = 11.2$, so that it may be compared directly to Fig. 5.25. As before, the signal to noise ratio is lower for $\hbar \approx 1$, and the tunneling signals are a bit weaker, even without added noise. In general, the system is much more sensitive to the effects of added noise than it was in the $\hbar = 2.08$ case. Roughly half as much noise is required to completely inhibit tunneling for the lower value of \hbar . This is reasonable because when we lower the effective value of the Planck constant the system tends towards classicality, where we are more sensitive to decoherence.

Chapter 6

Spatial focusing in the kicked rotor system

6.1 Overview

In all of the experiments that we have discussed thus far, we have studied momentum transfer to an atomic momentum distribution. In this chapter we describe our final set of experiments, where we instead study the evolution of the atomic *position* distribution. For various applications it is potentially useful to focus the atoms into a narrow region of space. Averbukh and Arvieu have recently proposed a new focusing method that is based upon an aperiodic kicked rotor potential [Averbukh01]. The primary context of their proposal is the rotational alignment of molecules in strong laser fields, which may also be described as a quantum rotor [Friedrich95]. However, the results are general and apply to any strongly driven quantum rotor. It should be possible to observe this effect by making simple modifications to our earlier kicked-rotor setup, and indeed we have recently observed the effect in our experiments.

Our experiments are conducted in the context of *atom lithography*, which is an emerging subfield of atom optics. Atom lithography is an atom-optic analogy to traditional optical lithography. Atoms incident on a surface are either directly or indirectly used to create structures on the surface. The indirect methods are the closest analogies of light lithography. In one realization, the energy released by the impact of metastable atoms is used to expose a resist layer on the substrate [Berggren95]. Fine structures in the position distribution of the incident atoms could in principle be created in the atomic distribution by selective optical pumping of atoms at certain locations [Chu96]. We are concerned with a direct method of atom lithography, where nanostructures are “grown” with the atoms that

are deposited onto the substrate. The tools of atom optics allow us to manipulate atomic center of mass motion with high precision in many ways. We can use upon the surface.

Atom lithography is an applied science and as such its goals are the development and improvement of certain technologies. One successful technique in direct deposition has been to focus the atoms with a standing wave of light. The dipole force on the atoms can act as a weak cylindrical lens that focuses atoms into an array of lines on a substrate [Timp92], as illustrated in Fig. 6.1. In the simplest realization, the lines are deposited with the periodicity of the standing wave, half a wavelength of the light. The feature size can typically be made much smaller than this spacing, producing features at the scale of tens of nanometers. This compares favorably with optical lithography techniques, which can produce features of order 100 nm. The spacing of the features can exhibit great regularity over macroscopic distances because it is defined by the regularity of the standing wave. An array deposited in this manner has several potential applications. The great uniformity of the structure can be exploited to form a precision distance scale with high resolution. The spacing of the features is determined by the wavelength of the light, which can be determined to high precision by using modern optical frequency metrology [Diddams00]. Such a scale may prove to be a useful secondary length standard. We can also imagine the array structure as a building block that may be translated to build complex microscopic structures. The inherently parallel nature of the system naturally lends itself to producing either large periodic structures or massive arrays of smaller devices.

Many of the experiments that have been performed to advance work in atom optics have been performed with hydrogen-like (alkali metal) atoms. Laser cooling and trapping techniques for these atoms are relatively simple compared with those of other types. However, these are not the sort of atoms that are useful for building nanostructures. Atomic focusing has also been demonstrated with certain technologically “useful” atoms such as chromium [McClelland93] and aluminum [McGowan95]. These metals are suitable for use in nanofabrication because of their conductivity, their adhesion to surfaces and low surface mobility, as well as their stability against exposure to air. Let us also note that the ability




Figure 6.1: Atom deposition with an atomic beam and a standing-wave lens.

to trap and cool a given species is not a prerequisite for atomic focusing. Many atomic species that are difficult to trap can be easily manipulated through the use of dipole forces. High atomic flux is a practical necessary for the construction of nanostructures over any macroscopic area. This can be easily achieved with a thermal atomic beam source, even for species that are not easily manipulated with laser cooling techniques. In the experiments that we will describe here, we continue to use cesium atoms. The results of this preliminary work can easily be applied for use with other elements.

Some progress has already been made towards the construction of more complex structures. A two-dimensional array of dots has been constructed with perpendicular standing-wave lens arrays [Gupta95]. The feature size of chromium lines has been made as small as 28 nm [Anderson99]. The distance between features can be made lower by using specialized lattice configurations [Gupta96] or Talbot-effect rephasing [Nowak97]. An array of chromium lines has served as a mask for the deposition of magnetic nanowires made of iron [Tulchinsky98]. A two-dimensional grid of magnetic dots created in this manner could potentially serve as a high-density magnetic storage device.

Focusing has also been demonstrated with a standing wave derived from a pulsed laser [Mützel00]. The description of this system is complicated by the fact that the adiabatic

approximation that was used to derive the strength of the dipole potential breaks down for very short pulses. On the other hand, pulsed lasers are available at wavelengths and powers that are sometimes difficult to obtain with continuous-wave (cw) lasers. This experiment was conducted in a regime where each atom was exposed to many (50) pulses on the way to the substrate, so that atomic motion in the standing wave was essentially the same as in the cw case.

While direct atomic lithography has been successful in many regards, there are several remaining problems. Besides the peaks, a broad background of atoms is typically deposited on the surface. In order to see the origins of these complications, let us consider the focusing process in slightly more detail. To be concrete, suppose that we have red-detuned light so that the atoms are focused towards the regions of maximum intensity. The potential can be approximated to be locally harmonic, with oscillation period T_{HO} . A perfectly collimated and monoenergetic sample of atoms in the harmonic oscillator potential will focus after a time $T_{\text{HO}}/4$. The final width of the distribution as it hits the surface is typically “diffraction limited” by the size of the ground state wavefunction in the harmonic potential. Unfortunately, the potential is not harmonic, and the initial sample is neither perfectly collimated nor monoenergetic. The sinusoidal nature of the potential leads to a spherical aberration that contributes to the background and the velocity spread of the distribution can lead to severe chromatic aberration. The transverse velocity spread of the incident beam (before the optical lattices) can be minimized by first cooling the beam with transverse (2D) optical molasses. The velocity spread of an atomic beam is typically low in the co-moving frame. An additional concern is that spontaneous emission events induced by the standing wave can lead to diffusion in momentum (and position to a lesser degree). Generally, this can be made negligible by an appropriate detuning choice. Besides effects that lead to the broad background, there are also sources of distortion that we must be concerned about. The stability of the substrate and the retroreflecting mirror must be maintained over the course of the deposition sequence. Because many experiments have used a standing wave that directly overlaps the substrate, the diffraction of the standing

wave off of the edge of the substrate can have significant effects [Anderson99].

Residual background effects pose a substantial problem for applications of these nanofabricated arrays. One solution has been to etch away a layer of the surface. This has been demonstrated with chromium atoms deposited on a silicon substrate to realize an array of nanowires [McClelland98]. Etching is very sensitive to the thickness of the layer that is being etched and it is not necessarily the case that the atomic beam is uniform over the entire region of deposition. If the substrate is etched enough to remove all of the background, the final feature width may not end up being particularly uniform. The etching process itself is not completely uniform, and may make this effect worse. Besides simple nonuniformities, the etching process tends to remove the silicon substrate that the chromium wires rest upon. In the most severe case, this can lead to the wire completely detaching from the substrate. In order to avoid these problems, we would like to minimize the amount of etching that is needed. The most natural way to do this is to develop new atom focusing methods that suppress the background.

The focusing mechanism that we have described so far is essentially that of coherent motion in a harmonic potential. The atoms tend to pass through the potential rather quickly, as they are moving towards the substrate at thermal velocities. If we were to begin with cold atoms instead of a thermal beam, the atoms can interact with the standing wave for a much longer period of time. Much as we loaded atoms into the standing wave for the SPASM sequence, it should then be possible to adiabatically focus the atoms into the potential minima [Khaykovich00]. In this regime, certain aberrations may become less severe and the background may be suppressed relative to the focusing that results from the coherent evolution alone. A combination of adiabatic and coherent focusing methods may also reduce the background from a thermal beam source. It is not yet clear if experiments will prove this to be an effective means of dealing with the background problem.

The proposal by Averbukh and Arvieu is a new approach to focusing that has the promise of reducing the background between the features. Let us suppose that an atomic beam that passes through a series of standing waves on the way towards the surface. We may

view this as a generalization of the process illustrated in Fig. 6.1. As the atoms pass through each standing wave, they momentarily experience the dipole potential of the standing wave. The time-dependent potential that the atoms experience is equivalent to that of a single kick in our kicked-rotor experiments. In view of this, we can test out our procedure “in place” by exposing atoms to a time-dependent series of pulses after being released from the MOT. We can further simplify the experiment by detecting the atomic focusing by some means other than depositing them on a substrate. These simplifications are adequate for a proof-of-principle experiment, but a more detailed study will necessitate actually focusing a beam onto a substrate. For such a study, the beam could be derived from a cold atomic source, in which case the overall procedure would be very close to that of our experiments.

Let us now come to the focusing method itself. With a single kick, the focusing is not very different from the coherent focusing that has been demonstrated in previous experiments. The potential is again to some degree harmonic and gives each atom a kick in momentum towards the potential minimum at $x = 0$. Some time later at time T_f the atoms reach their focus. This not a terribly sharp focus because it is affected by the aberrations that we have previously described. It turns out that if we wait slightly longer than this simple focusing time, the atoms that were not otherwise efficiently focused continue to move towards the potential minimum. The result is that the maximum compression of the position distribution occurs at time T_c , which is slightly after T_f . Exactly what we mean by “compression” here is a technical matter, but we can define some figure of merit f_c that describes how localized in position the distribution is. The time T_c is the time at which the compression figure of merit is optimized. If we apply a second kick at time T_c , we can further increase the spatial compression. With a series of several appropriately spaced kicks, we should be able to achieve high spatial compression while suppressing the background. The time scale for compression after each subsequent kick decreases monotonically as the typical velocity of the distribution increases. Ultimately, the compression that can be achieved in our experiments is limited by the finite (nonzero) duration of the pulses that we use. We employ a series of kicks of constant duration and amplitude that become closer to each

other in time. During each kick, the atoms are exposed briefly to the one-dimensional optical lattice. The average intensity of the lattice increases gradually in time as the time between kicks becomes shorter. We then note that in some sense, this is similar to the idea of adiabatic focusing.

6.2 The Aperiodic Kicked Rotor

A more formal description of the experiment begins with the Hamiltonian for the aperiodic kicked rotor system. The Hamiltonian for the periodically kicked rotor is given by Eq. 1.47. We now relax the condition that the kicks are equally spaced in time. The Hamiltonian is then given by

$$\mathcal{H}(x, p, t) = \frac{p^2}{2m} + V_0 \cos(2k_L x) \sum_n F(t - t_n), \quad (6.1)$$

where the kicks occur at times $\{t_n\}$. As in the case of motion in the pendulum potential (§1.4.4), there is not a natural time scale for motion in this system in the absence of external driving. We can then pick a set of scaled units for this system where we choose $\hbar = 1$ for the sake of convenience. That choice fixes the time unit of the system to be $T_u = 1/8\omega_r$. The remaining transformations are in parallel with those that we used for the periodic kicked rotor (§1.4.3). We define the pulse integral as $\eta = 8\omega_r \int_{-\infty}^{\infty} F(t) dt$, such that η is proportional to the pulse width t_p . Furthermore we take

$$\begin{aligned} x' &= 2k_L x, \\ t' &= t/T_u = 8\omega_r t, \\ p' &= p\hbar/2\hbar k_L = p/2\hbar k_L, \\ f(t') &= F(t)/\eta, \\ K &= (\hbar/2\hbar) \eta T_u V_0 = (\eta/8\omega_r \hbar) V_0, \\ \mathcal{H}' &= (\hbar/2\hbar) T_u \mathcal{H} = \mathcal{H}/8\omega_r \hbar, \end{aligned} \quad (6.2)$$

and then drop the primes. Under these transformations Eq. 6.1 becomes

$$\mathcal{H}(x, p, t) = \frac{p^2}{2} + K \cos x \sum_n f(t - t_n), \quad (6.3)$$

where we identify K as the stochasticity parameter. In the limit where the pulse time goes to zero at constant K , we recover a generalized version of the δ -kicked rotor. As in the

case of the periodically kicked rotor, we can integrate the classical equations of motion over one kicking period to obtain

$$\begin{aligned} p_{n+1} &= p_n + K \sin x_n \\ x_{n+1} &= x_n + \Delta t_n p_{n+1} \end{aligned} \quad (6.4)$$

a variation of the standard map. In writing down the mapping, we have defined $\Delta t_n \equiv t_{n+1} - t_n$ to be the time between subsequent kicks.

The focusing time after a single kick is $T_f = 1/K$, for atoms close to the potential minimum. In order to determine the time of maximum compression, we must define a suitable figure of merit. One possible measure is the orientation factor $f_c = \langle \cos x \rangle + 1$, where we are averaging over the entire atomic distribution. For a distribution that is completely uniform across a period of the standing wave, $f_c = 1$. If the distribution is completely localized at the potential minimum, we would have $f_c = 0$. The exact choice of the figure of merit is not important, so long as it is one that properly quantifies the degree of compression. We can imagine that f_c and some other appropriate figure of merit might be minimized at times very slightly offset from each other. The next kick in the sequence would then occur at slightly different times with the two measures, and it is likely that one of the two sequences would more efficiently focus the atoms. While it is not obvious that any particular figure of merit is optimal, the variations are presumably small and can be made up with a few additional kicks. The time T_c of maximum compression occurs after T_f , and we can further distinguish the two by noting that the T_f is the time when the position distribution is maximally peaked, but the distribution is more compressed overall at time T_c . A comparison of these features is illustrated with simulations in Fig. 6.2.

The long-time behavior after a single kick is not a matter of great importance because we kick the distribution at the *first* minimum of f_c . Nevertheless, there are some important features in the longer time dynamics that affect the procedures that ultimately affect the simulations that we employ. As we have discussed in other contexts, one fundamental difference between our system and a true quantum rotor is that there is more than one possible momentum ladder in our experiment. This has serious implications for the observation of some purely quantum phenomena including quantum resonances (§1.3.3)




Figure 6.2: Calculation of atomic focusing after a single kick. These curves are the result of a simple quantum simulation for a single δ -function kick of strength $K = 5$. We first plot the total amplitude contrast of the atomic spatial distribution (blue, dashed). This function is the position distribution contrast $-1 \cdot (\text{Max}(I(x)) - \text{Min}(I(x)))$, where $I(x)$ is the atomic position distribution (intensity). We have normalized the contrast figure to the amplitude of the other curve for ease of comparison. This contrast functional is minimized when the atomic position distribution is most sharply peaked. For a fully harmonic potential we would expect the focusing time T_f to be $1/K = 0.20$, but it turns out that the focusing time here is closer to 0.26. We also plot $f_c = 1 + \langle \cos x \rangle$ (red, solid). This functional is minimized at $T_c \approx 0.357$, when $f_c \approx 0.442$. The most important feature to note here is that this minimum occurs at a substantially later time than the minimum in the other curve. The initial condition for the simulations is close to that of our experiments which use a lattice-cooled distribution. The distribution is Gaussian in momentum with width $\sigma_p = 0.8$ and is uniform in position over the scale of the standing wave.

and dynamical tunneling (§5.2.2). In the present case, the oscillations in f_c repeat with period 4π for the case of a momentum ladder centered at $p = 0$, which has symmetry under momentum reversal. For the case of a continuous set of momentum ladders, the behavior up until time $t = \pi$ is indistinguishable from the single-rotor case, but it ceases to display




Figure 6.3: Comparison of particle and rotor behavior after one kick. In the case of a true quantum rotor (dashed) where there is only one momentum ladder, the value of f_c undergoes oscillations that repeat with period $t = 4\pi$. If we instead begin with particles distributed across a set of momentum ladders (solid), the long-time oscillations are not observed. These simulations are conducted with $K = 5$ and δ -function kicks. The difference between the two cases is that the momentum grid for the rotor case has unit spacing, while the other case has 32 divisions within the unit. The behavior of the two cases does not deviate until well after the first minimum of f_c , which justifies the use of the coarse grid in the other simulations presented here.

oscillations after this time. This is illustrated in Fig. 6.3. Since we are interested in the dynamics at shorter times than this (when the two cases agree), it is usually sufficient for us to perform simulations assuming a single momentum ladder. As a side note, Averbukh and Arvieu suggested that it might be possible to wait a full quantum resonance period before applying a kick at the subsequent minimum of f_c . This is not by any means *necessary*, but under certain circumstances it might make aspects of the experiment easier. For the reasons that we have just described, waiting for this period is not feasible unless the initial condition has a single momentum ladder. Although for our CAT experiments we used stimulated Raman velocity selection, this is a rather inefficient way to select a momentum ladder. A




Figure 6.4: Calculation of atomic focusing after a multiple kicks: f_c . This simulation shows the evolution f_c (solid), as calculated in Fig. 6.2. Again we have $K = 5$ and we are assuming that the kicks are δ -functions. Instead of simply tracking the evolution after a single kick, we now apply a new kick each time that f_c reaches a local minimum (circles). We follow the evolution of f_c for the first 30 kicks. As the kick sequence proceeds, the kicks become closer together. Eventually, the circles that mark the kick times overlap. If we suppose that the finite duration of the kick were represented by the width of the circles, it becomes clear that the pulse width may become a limiting factor in the experiment.

better method would be to begin the experiment with an “atom laser” beam derived from a BEC, which would be extremely cold in the transverse directions. Unfortunately, atom lasers are not yet available with any technologically useful atomic species.

It is with the application of additional kicks that the background becomes efficiently suppressed. Each subsequent kick is applied at the time T_c that minimizes the orientation factor f_c , a time Δt_n after the previous kick. When we follow this procedure, the value of f_c decreases steadily with additional kicks. We illustrate this behavior with simulations in Figs. 6.4, 6.5 and 6.6. Only a few kicks are necessary to substantially reduce the overall




Figure 6.5: Calculation of atomic focusing after a multiple kicks: position distribution evolution. This simulation shows the evolution the position distribution just after each kick for the simulation shown in Fig. 6.4. The position distribution is represented by the color, where the most intense regions appear white. The leftmost slice (after zero kicks) is completely uniform, representing the initial condition. The horizontal axis represents the kick number, not the time— it is important to note that the time between kicks is much shorter for the later distributions. While the distribution becomes more compressed as the background is suppressed, it retains complex structures that result from repeated diffraction off of the standing wave. Some of the individual distributions from this figure are shown in Fig. 6.6.

background. As additional kicks are applied, the remaining width of the distribution gradually decreases. The position distribution is not smooth, but is composed of sharp maxima that are reminiscent of caustics encountered in rainbow scattering [Averbukh01]. Although the time between kicks and the corresponding changes in f_c are erratic for the first few kicks, the behavior becomes smooth on a longer time scale. For true δ -function pulses, there is no clear limit to this process.




Figure 6.6: Calculation of atomic focusing after multiple kicks: individual distributions. Several of the individual position-space distributions from Fig. 6.5 are shown here in detail. The distribution is shown in (a) after 0 kicks (green dashed), 1 kick (blue), 2 kicks (red), and 4 kicks (black). In (b) the distribution is shown after 5 kicks (green dashed), 10 kicks (blue), 15 kicks (red), and 20 kicks (black). The most important feature that we can see in this plot is that the background becomes suppressed after only a few kicks.

For pulses of finite duration, we can expect that the process will break down when the pulse duration interferes with the free evolution that is necessary between subsequent kicks. One partial solution is to wait one quantum resonance period (plus the desired time between kicks) before applying the next kick. As we have mentioned, this is generally not a practical solution. We must also be concerned about atomic motion during the kick itself. During each kick the optical lattice has a well depth given by $\alpha = K/t_p$, supposing that the pulses are square. We can then estimate the best focusing that we can obtain with pulses of duration t_p . We are interested in the behavior of a highly localized wave packet that is located near the potential minimum of the lattice. In this regime the potential is approximately harmonic and the smallest distribution that we can expect is the minimum-uncertainty ground state for this well depth. The ground state is Gaussian, with width $\sigma_x^2 = 1/2\sqrt{\alpha} = \frac{1}{2}\sqrt{t_p/K}$. For a Gaussian distribution with this width, the orientation is given approximately by $f_c = \frac{1}{4}\sqrt{t_p/K}$. This is an asymptotic expression, and we have implied that the state is one of minimum uncertainty, which we would not expect from the multiply peaked position distributions that we have seen thus far. In general it is necessary to perform numerical simulations in order to determine the limiting values. Simulations confirm that this analytic estimate for f_c is indeed a lower bound.

6.3 Experimental Methods

The procedures and data acquisition sequences for our atomic focusing experiments have a high degree of overlap with the other experiments that we have described. The overall sequence is very close to that which we described in Chapter 2. Beyond this, we apply the 3D lattice cooling sequence (§4.2) to further reduce the initial momentum spread of the atomic distribution. These experiments were performed after the development of the state preparation procedures, but we did not utilize any of the more advanced techniques here. One detail that we should note at this point is that the electro-optic phase modulator (EOM) is still placed before the retroreflector on the interaction beam. This allows us to shift the phase of the standing wave, which is important for some of our detection methods

as we will see shortly.

The most important change from the earlier experiments is that we are now interested in the *position* distribution rather than the momentum distribution. It is relatively easy for us to measure the momentum, since a simple ballistic expansion procedure converts the momentum distribution into a macroscopic position distribution that we can image. There is not an obvious and simple method for us to determine the shape or size of the position distribution on the scale of one period of the standing wave. Our strategy then is to develop some means of extracting information about the microscopic position distribution through one of the quantities that we can measure. We will now describe several potential methods for measuring the microscopic position distribution. After developing these methods, we will describe some of our calibration procedures and the overall experimental sequence for these experiments.

6.3.1 Near-resonant detection

A standard technique that has been used to determine the average position of atoms with respect to a standing-wave axis is to induce fluorescence with the standing wave. The rate of spontaneous emission events depends sensitively upon the intensity of the light, which has well defined maxima and minima for a standing wave. It is often possible to interact only with atoms in a relatively small spatial region. By changing the phase of the standing wave, the region of interest can be swept across an entire period of the standing wave to map out the distribution. This is an example of a “virtual slit” [Chu96] detection method. Two possible variations are (a) to count the fluorescence directly or (b) to promote atoms at specific locations to a different internal state.

For either of these methods, we first optically pump our cold cesium atoms into the $F_g = 4$ ground state. We then apply the focusing-kick sequence to the atoms using a standing wave directed along the x -axis. (Generally speaking, we try to work in a regime where the far-detuned standing wave does not substantially affect the internal state of the atomic sample.) We next expose the atoms to a brief pulse of light from a standing wave

(along the same axis) that is sufficiently near resonance to induce spontaneous emission. Each spontaneous emission results in fluorescence that can be directly observed and also has some probability of moving the atom to the other ($F_g = 3$) ground state. These two processes are most likely to occur at the antinodes of the standing wave, where the light is most intense, and unlikely to occur at the nodes. We can then change the relative position of this wave and the atomic sample by applying a single phase shift to the standing wave after the kicking sequence but before the detection pulse. By performing repeated measurements at different phases, we then determine the relative atomic population at different positions along the x -axis. As in other cases, the detection pulse must be short enough that the atoms do not have time to move (on the length scale of the standing wave period) during the pulse.

Both the direct and indirect methods are reasonable approaches, however it is not practical for us to directly count the fluorescence at this time. We do not presently have a high-sensitivity device such as a photomultiplier tube or an avalanche photodiode (APD) to detect fluorescence from the MOT, although we are in the process of adding an APD system to our setup. Furthermore, it does not appear to be feasible to use our CCD camera to measure such a small amount of fluorescence.

The second near-resonant method is somewhat similar to the stimulated Raman tagging technique (§4.3) that we have used in other contexts. However, we are now interested in tagging a set of atoms at a specific position, rather than momentum. We again begin with cold atoms in the $F_g = 4$ ground state. After the kicking sequence, we again apply a pulse of near-resonant light at some specified phase. Only the atoms that undergo spontaneous emission events are transferred to the $F_g = 3$ state. We then push away all of the atoms that remain in the $F_g = 4$ ground state (§4.3.3). We then image the remaining atoms using our standard freezing molasses technique to count the number of $F_g = 3$ atoms. The number of atoms transferred to this state describes the overlap between the atomic position distribution and the standing wave intensity profile at the given detection phase.

We attempted to apply this spatial tagging method in our experiment, but without success. Our first efforts used a near-resonant standing wave derived from the MOT mo-

lasses beams. The beam was the resonant $F_g = 4 \rightarrow F_e = 4$ beam that we use elsewhere in our experiment for optically pumping atoms to the $F_g = 4, m_F = 0$ state prior to stimulated Raman tagging. We combined this light with the interaction beam shortly before the chamber by using a polarizing beamsplitter cube. The measurements that we made with this setup were extremely noisy. A simple first test of a detection method is to see if it can give an acceptable “zero” signal. In this context, the zero signal results from the detection procedure without the focusing-kick sequence, where we expect to see a uniform spatial distribution. As it turns out, the noise level was larger than the average signal level of this supposedly straight line. One possible cause is the linewidth of the DBR laser (from which this beam is derived), which can affect the spontaneous emission rate in unpredictable ways. We also tried detuning the beam several megahertz off of resonance by driving the AOM controlling that beam at a different frequency, but it did not solve the problem.

It is also possible to induce spontaneous emission events by using the interaction beam itself as a source of quasi-resonant light. One advantage of this method is that there is no frequency difference between the kicking and detection beams, so the absolute phase shift between the two beams is easily determined. In the case that spontaneous emission is induced by the $F_g = 4 \rightarrow F_e = 4$ beam, the 18 GHz or so of frequency difference between the two beams leads to a phase shift of order 30λ over the 0.5 m (or so) distance between the retroreflecting mirror and the atoms. Another advantage is the simplicity of the setup—it is not necessary to introduce any additional beams into the chamber. On the other hand, if we wish to actually induce an appreciable number of spontaneous emission events, it is necessary that we detune in closer to resonance. Our implementation of lattice cooling has not been found to work at detunings closer than 12 GHz red of resonance. At 12 GHz, we did not see a substantial population in the $F_g = 3$ state, so we tried detuning in closer, where we could not use lattice cooling. Although we tried detunings as low as 1.5 GHz, we were not able to obtain a zero-signal measurement that was substantially less noisy than that obtained with the $F_g = 4 \rightarrow F_e = 4$ beam.

6.3.2 Asymmetric kicking measurements

In the absence of reliable methods for determining the position distribution with near-resonant light, we were forced to develop non-resonant methods for determining the spatial distribution. The simplest non-resonant detection method that we have employed involves a single standing-wave δ kick that we use to probe the microscopic spatial distribution of the atoms. At a variable time after the end of the squeezing-kick sequence, we apply a single, strong pulse. We used the same retroreflected beam that delivers the focusing kicks, but at 80% of the available laser intensity for 600 ns. As with the resonant detection methods, we apply a phase shift between the end of the focusing-kick sequence and the beginning of the detection pulse. Because of the nonzero response time of the EOM that we use to shift the phase, the phase shift (in all cases) begins 1 μ s before the detection pulse begins.

The detection pulse transfers a different amount of momentum to the atoms depending on their position with respect to the standing wave. An atom that is at the node or antinode of the detection beam does not feel any force from the standing wave, whereas one directly between them will be kicked strongly to one side or the other (depending on the phase of the standing wave). Now, if we consider an atomic distribution that is strongly localized in position, we expect to see the same behavior: a strong dependence of the final momentum distribution on the phase of the detection pulse. In contrast, if the momentum distribution is uniform on the scale of a standing-wave period, the final distribution will be independent of the detection phase. In practice, this gives us a way to measure how localized the distribution is in position. If we kick the atoms with a standing wave with several different phase shifts, a spatially localized distribution acquires an average positive or negative momentum depending on the phase. The contrast Δ_p in $\langle p \rangle$, which we define as $\Delta_p \equiv \langle p \rangle_- - \langle p \rangle_+$, is thus a useful figure of merit for the degree of phase-space localization. The value $\langle p \rangle_+$ is the expectation value of p when we kick the distribution with a phase shift of $\pi/2$ and $\langle p \rangle_-$ is the value when we shift the phase by $3\pi/2$ (or $-\pi/2$). Note that the sign of Δ_p reverses along with the direction that the atoms are kicked by the standing wave. In the case that we do not apply a phase shift, we typically observe two cleanly separated




Figure 6.7: Effects of finite detection pulse duration. These curves are similar to those that appear in Fig. 6.2, and show the evolution of the position distribution functionals f_c (dashed) and f_p (solid) after a single 300 ns focusing kick where $K = 3.2$. The value of f_p is that which we would expect with the detection pulse duration (600 ns) and kicking strength ($K_{\text{det}} = 7.33$) that we use in the experiment. In the limit that the pulse width t_p goes to zero, this curve converges to f_c . The minimum of the f_p curve appears slightly before that of f_c . However, when f_p is minimized, f_c is also very close to its minimum. As we have mentioned, the overall focusing procedure is not very sensitive to the particular choice of the figure of merit. The amplitude of the two measures is not exactly the same, but they are very close.

peaks that appear at about $\pm 5 \cdot 2\hbar k_L$ after the detection pulse and free drift time. With a phase shift of $\pi/2$, a fully localized distribution will entirely be kicked to about $5 \cdot 2\hbar k_L$. We then expect the contrast figure to be of order $\Delta_p = 10 \cdot 2\hbar k_L$, for an optimally localized distribution and zero for a uniform distribution. It is practical for us to use the figure of merit $f_p \equiv 1 - (1/2K_{\text{det}})\Delta_p$, where K_{det} is the strength of the detection pulse.

Let us consider the effects of this measurement in slightly more detail. Suppose that we begin with an initial distribution described by the state vector $|\psi_0\rangle$. Often it is

convenient to work with the corresponding position-space wave function $\psi_0(x) = \langle x|\psi_0\rangle$. The single detection pulse in our experiment is approximately a δ -function kick, which we can describe with a quantum mechanical operator (§1.3.1). After the kick, the wave function is given by

$$\psi(x) = U_0(K, \phi)\psi_0(x) = \exp\left(-\frac{i}{\hbar}K \cos(x + \phi)\right)\psi_0(x), \quad (6.5)$$

where $U_0(K, \phi)$ is a single kick operator with strength K and phase ϕ . After the kick, we measure the average momentum with two initial phases to find

$$\Delta_p = \langle U_0(K, -\pi/2)\psi_0|p|U_0(K, -\pi/2)\psi_0\rangle - \langle U_0(K, \pi/2)\psi_0|p|U_0(K, \pi/2)\psi_0\rangle. \quad (6.6)$$

We can now calculate $\langle p\rangle_s = \langle \psi_s|p|\psi_s\rangle = \langle U_0(K, s \cdot \pi/2)\psi_0|p|U_0(K, s \cdot \pi/2)\psi_0\rangle$, where s is one of either ± 1 . We find that

$$\begin{aligned} \langle p\rangle_s &= \int \psi_s^*(x) \frac{\hbar}{i} \frac{\partial}{\partial x} \psi_s(x) dx \\ &= \int \psi_0^*(x) \exp\left(\frac{iK}{\hbar} \cos(x + \frac{s\pi}{2})\right) \frac{\hbar}{i} \frac{\partial}{\partial x} \exp\left(-\frac{iK}{\hbar} \cos(x + \frac{s\pi}{2})\right) \psi_0(x) dx \\ &= \int \psi_0^*(x) U_0(-K, s\pi/2) U_0(K, s\pi/2) \frac{\hbar}{i} \left(\frac{\partial}{\partial x} + \left(\frac{iK}{\hbar}\right) s \cos(x)\right) \psi_0(x) dx \\ &= \int \psi_0^*(x) \left(\frac{\hbar}{i} \frac{\partial}{\partial x} + sK \cos(x)\right) \psi_0(x) dx \\ &= \langle p\rangle + sK \langle \cos x\rangle, \end{aligned} \quad (6.7)$$

where $\langle p\rangle = \langle \psi_0|p|\psi_0\rangle$ is the average momentum of the initial (unkicked) state. Finally we obtain a value for the total momentum contrast,

$$\Delta_p = \langle p\rangle_- - \langle p\rangle_+ = -2K \langle \cos x\rangle. \quad (6.8)$$

This result determines our definition of the function $f_p = 1 - (1/2K)\Delta_p$, which turns out to be an equivalent measure to f_c when $K = K_{\text{det}}$ is the strength of the detection pulse. This measurement technique is subject to many of the systematic effects that we encountered in the context of the earlier kicked rotor experiments (§3.6.2). The final extent of the distribution in position is slightly broadened by the initial spatial extent of the MOT. More importantly, the nonzero duration of the detection pulse leads to a slight reduction in the effective value of K for the pulse. The mechanism for the reduction is atomic motion in

the optical lattice during the pulse. A more severe consequence is that the *time* that the function is minimized can change, as illustrated in Fig. 6.7. We have used these momentum-contrast measurements in our experiments to calibrate the time between subsequent kicks, so this can affect the cumulative efficiency of our focusing procedure. As before, the most appropriate method of accounting for these effects is to directly incorporate the pulse width into simple quantum simulations. The systematic effects can be minimized by using pulses that are as strong and brief as possible.

6.3.3 Pendulum oscillation measurements

A second non-resonant detection method is to observe the behavior of the wave packet after it is projected into a pendulum phase space. After the final kick, we allow the system to evolve until the time when the maximum squeezing occurs, as determined by the momentum-contrast measurements. This is the point in time when the next kick *would* occur if there were to be additional focusing kicks. At that time, we project the distribution into the pendulum phase space at a specified amplitude for a specified amount of time. After this period of motion in the lattice, we measure the momentum distribution in the usual way. We repeat this procedure many times to trace out the two-dimensional parameter space of well depth and time. At each value of the well depth we observe the time evolution of the momentum distribution. The sampling rate in time is high enough that we can observe the atomic wave packet oscillating repeatedly in the potential. We can imagine that if the initial atomic wave packet were closely matched to the ground state of the optical lattice, we would project atoms mostly into that state. The potential is approximately harmonic, so the ground state is a minimum-uncertainty Gaussian stationary state. By changing the well depth, we can change the phase-space aspect ratio of this state. Optimally, we would like to see some value of the well depth at which the amplitude of the oscillations is minimized. This would indicate that we have maximized the overlap between our focused distribution and the harmonic-oscillator ground state. The width of the ground state is given by $\sigma_x^2 = 1/2\sqrt{\alpha}$, so that as long as we know α , we have a direct means of measuring the spatial size of the wave packet. Beyond simply measuring the feature size, we can hope to get some indication of

the residual background with this measurement. If most of the atoms are projected into the ground state at some well depth, any oscillations are likely to be dominated by the residual background of unfocused atoms.

It is necessary to define some figure of merit that quantifies the amplitude of the oscillations. One possibility is to fit each momentum distribution to a Gaussian function and to extract the width. Unfortunately, the distributions are not necessarily Gaussian in nature and this fitting procedure fails. A more natural measure is the ensemble energy $E = \langle p^2 \rangle / 2 = \int_{-\infty}^{\infty} p^2 / 2 f(p) dp$, where $f(p)$ is the momentum probability distribution. As we have discussed in the context of earlier kicked rotor experiments, we cannot determine the energy with great accuracy or precision. This is because the energy is weighted heavily by the behavior at large momentum values, where the signal to noise ratio is the lowest and the nonlinearities in the detection system become important. Great computational effort may be required to directly account for these issues, especially when the system is not fully described by classical dynamics. A measure of the width of the distribution that is less sensitive to the behavior at high momentum values is the Süßmann measure [Süßmann97], brought to our attention by Prof. Schleich [Schleich01]. The Süßmann measure is defined as

$$\delta_{Sp} = \frac{1}{\langle f(p) \rangle} = \frac{1}{\int_{-\infty}^{\infty} f^2(p) dp}, \quad (6.9)$$

the inverse of the area under the square of the distribution function. This function emphasizes the distribution in the region where it is the most intense and the signal to noise ratio is the highest, near $p = 0$. As a result, the evolution of δ_{Sp} is typically a smooth curve that is much less noisy than the energy measurements. One nice feature of the Süßmann measure is that it is a rather general functional that can tell us the width of distributions of different shapes. In the experiment, the value of δ_{Sp} is computed at each step in the time evolution. We would then like to determine the amplitude of the oscillations by examining the $\delta_{Sp}(t)$ curve. Since this curve is not necessarily a simple function of time, it is not practical to fit it to find the amplitude of oscillation. However, the standard deviation σ_S of $\delta_{Sp}(t)$ is a practical measure of the amplitude.

6.3.4 Free expansion measurements

The measurement procedures that we have described so far are ones that measure quantities related to the atomic position distribution. It is also possible that the momentum distribution could be of interest. We expect heating to result from the sequence of kicks, and it is useful to characterize exactly how much heating there is. Furthermore, we wish to fully understand the pendulum oscillation measurements. For this, it is useful to have some idea of the shape of the momentum distribution. The momentum distribution for these experiments can be measured in the usual way, by performing a temperature measurement after the kicking sequence. No additional “detection pulse” is necessary in this case.

A variation on this method can be used to provide an additional measurement of the position distribution. After the kicking sequence, we project the distribution into the pendulum potential as we did in the pendulum oscillation measurements. In this case, we pick the deepest well depth that is available, so that the potential most closely approximates that of a harmonic oscillator. Furthermore, we only allow the atoms to evolve in the lattice for $T_{\text{HO}}/4$, one quarter of the harmonic oscillator period. This evolution in the potential converts the atomic position distribution into a momentum distribution, and the momentum distribution into a position distribution. If we perform a temperature measurement after this detection pulse, we have an additional determination of the position distribution.

It is important to point out that the anharmonicity of pendulum potential can still have significant effects, as illustrated in Figs. 6.8 and 6.9. In these figures we study classical evolution in the pendulum potential with both a Monte Carlo simulation and an analytic solution. In both cases, we begin with a distribution that is comparatively narrow in momentum and delocalized in position. After a time $T_{\text{HO}}/4$, we plot the momentum distribution. The final distribution shows peaks at the edges with a width that is determined by the initial width of the momentum distribution. If the potential were purely harmonic, we would not expect to see these peaks. If the initial condition is a delta function, the momentum distribution exhibits singularities at the peaks. For any finite distribution width the peaks will be finite. Likewise, the finite extent of a quantum wave packet is sufficient




Figure 6.8: Pendulum Rotation: classical Monte Carlo simulations. In (a) we follow the motion of 10^4 particles through the phase space. The pendulum well depth is $\alpha_p = 106$. The sample of particles (red) is uniformly distributed in position and has momentum width $\sigma_p = 0.8$. After time $T_{\text{HO}}/4$, the distribution (blue) has rotated significantly in the phase space. The central part of the distribution (where the distribution is most harmonic) has rotated by 90 degrees. Part (b) shows histograms of the momentum distributions in part (a). These two curves are normalized such that the maximum intensity of the initial distribution (red) is unity.

to replace the singularities with a sharply peaked wave function. We can also convolve this distribution with a Gaussian function that represents the initial momentum width of the




Figure 6.9: Pendulum Rotation: classical analytic solution. The classical equations of motion for the pendulum yield a solution in terms of Jacobi elliptic functions, which is plotted here (a) for a distribution that begins (red) with $p = \delta(0)$, uniformly distributed in x . After evolving for time $T_{\text{HO}}/4$, the distribution (blue) follows a curve like that shown in Fig. 6.8. The final momentum probability distribution (b) is difficult to plot accurately because it is so sharply peaked— it is apparent from the curve in part (a) that these peaks are singularities. The distribution is plotted with a large number (10^3) of momentum bins.

distribution to produce a curve very similar to that which results from the Monte Carlo simulation. The resulting *rainbow function* is the final momentum distribution that we

expect when the initial distribution is uniform in position. In the experiment then, we wish to examine the total momentum distribution after a quarter-period of oscillation in the pendulum potential. If there is a component of the distribution that is described by the rainbow function, then the amplitude of that component is proportional to the residual (uniform) background of the distribution.

6.3.5 Well depth calibration

In many other contexts, we have noted that it can be difficult to calibrate the absolute well depth of our optical lattice. In the earlier generation of kicked-rotor experiments, we used the locations of the peaks in the $D(K)$ curve to perform the absolute calibration (§3.3). It turns out that we cannot perform this particular calibration here because we do not have enough laser power to observe multiple peaks at the larger detuning that we are using for these experiments. For the chaos-assisted tunneling experiments, we placed a minimum-uncertainty wavepacket into the pendulum potential and observed the oscillations to calibrate the well depth. This is practical in that case because the atoms all move in a bunch together through the phase space. It is relatively easy to extract the frequency of oscillation. When our only state preparation is lattice cooling, this initial state is wide enough that oscillation in the well is no longer a simple problem. Furthermore, the state preparation procedures work best at higher detunings, and it is not practical to use those methods to calibrate our well depth for this experiment.

We can always use our setup to generate a simple potential, where the dynamics depend upon the lattice well depth. If the overall system is simple enough, we can hope to computationally model the dynamics and fit the data to find the well depth. Such a tractable problem in the context of these experiments is the momentum contrast measurement after a single focusing kick. When performing the experiments we define the instantaneous laser intensity to a fixed reference point that we refer to as the available laser intensity. The purpose of the calibration procedures is to find the absolute value of this intensity, which is not initially known. In these measurements we exposed the atoms to a 300 ns focusing

kick, at 10%, 25%, or 75% of the available laser intensity. This is followed by some time delay and then a 600 ns detection pulse at 80% intensity. The pulse widths and their *relative* intensities are well known. By incorporating these parameters into the simulation we can reproduce the curves that are given by the experiment, and iteratively refine the scaling factor that determines the absolute well depth. The simulations account for the finite duration of the pulses, and only simulate the system up until time $t = \pi$ so that the system is effectively described by a single momentum ladder. The result of these simulations determines the 100% intensity level during the focusing kicks to correspond to $K = 4.26 \pm 4\%$. The uncertainty is the purely statistical error associated with the fitting procedure, and does not account for any of the other factors in determining K . During the data runs we constrained the laser power to stay within $\pm 2.5\%$ using the VPL system (§2.3.3). Beyond the drifts associated with time variation of the lattice intensity, we must be concerned with other mechanisms that broaden the value of K . As in our previous experiments with the kicked rotor, the finite pulse duration leads to a reduction in the effective value of K as a function of momentum (§3.2). This particular effect is accounted for by the finite pulses used in the simulation. Most of the other systematic effects are not accounted for, for example the spread in the effective value of K due to the transverse profile of the interaction beam. A conservative estimate of the maximum kick strength is $K = 4.26 \pm 10\%$. Based upon this calibration and the known area of the detection pulse, we have determined the detection pulse kicking strength to be $K_{\text{det}} = 7.33 \pm 10\%$.

6.3.6 Experimental sequence

As we have discussed, these experiments are very similar to the earlier kicked rotor experiments, with the exceptions of the lattice cooling and the detection methods. The Ti:sapphire laser that generates both the 3D lattice cooling beams and the interaction beam was detuned 18 GHz to the red of the cycling transition for all of the experiments described here. Generally speaking, the signal-to-noise ratio of these experiments is fairly high because our initial distribution is fairly cold and the lattice-cooled distribution contains a fairly large number of atoms, especially compared to those experiments with the stimulated-Raman velocity

selection. We begin each data set with two background images on the CCD camera.

We begin by loading the MOT for 5 s. The maximum number of atoms, based upon atomic fluorescence counting with the CCD camera is estimated to be 10^6 . We use 3D lattice cooling with the same parameters (except for the Ti:sapphire laser detuning) as we described earlier (§4.2). The typical temperature after lattice cooling (without optical pumping) is roughly 500 nK. This temperature is slightly higher than it was in the earlier case because of the difference in detuning. Typical residual magnetic fields at the end of lattice turn-off time are of order 70 mG and do not change substantially during the course of the experiments. We next perform optical pumping to $F_g = 4$ state by turning on the repumping light for 50 μ s. The typical temperature after lattice cooling and full optical pumping sequence is roughly 625 nK, corresponding to $\sigma_p = 1.8 \hbar k_L$. Without lattice cooling, the final temperature is roughly 18 μ K, corresponding to $\sigma_p = 9.5 \hbar k_L$.

The kick sequence consists of a series of 300 ns pulses of constant amplitude. The overall pulse sequence may last as long as 50 μ s. In general, the pulse sequence is not periodic. After the focusing kicks, we begin the detection sequence with one of the methods that we have described. The essential methods are the momentum contrast, the pendulum oscillations, and the free expansion measurements. In the case of the momentum contrast method, we wait some (variable) amount of time between the end of the kicking sequence and the detection pulse. In order to ensure even response of the EOM that shifts the phase of the interaction beam, the EOM begins shifting the phase of the light to the desired phase at a time 1 μ s before the beginning of the detection pulse. For the other detection methods, there is no phase shift or delay time required. The ballistic expansion period is 25 ms, and the CCD exposure time in the freezing molasses is 15 ms.

6.4 Data and Results

The focusing procedure employs a series of kicks of constant amplitude that gradually become closer in time. The limiting factor for how close the pulses can come is the response time of the EOM, for which a constant 1 μ s of slew time is allowed. This limitation is




Figure 6.10: Experimentally measured time evolution of the momentum contrast figure of merit, f_p after 1, 2, 4, 6, 8, and 10 kicks at $K = 0.42$. The time (horizontal) axis is calibrated in microseconds.

specific to the momentum-contrast detection method that we use. If we used some other detection method, it would be possible to let the pulses come much closer together. An optimal kicking sequence with finite pulses would proceed until the focusing parameter f_c saturates to its minimum value for the given pulse width. Intuitively, we might expect this saturation to occur when the time between subsequent pulses becomes comparable the duration of the pulses. However, the slew time in the momentum-contrast measurement is much longer than the pulse duration and so it becomes the limiting factor. The number of kicks that is required to reduce the focusing time to less than $1 \mu\text{s}$ depends upon the strength of the kicks. We studied focusing in this system with the same three amplitude values that we used for the well depth calibration, 10%, 25%, and 75% of the available laser intensity. These values correspond to K of 0.43, 1.06, and 3.20, respectively.

For each of value of K , we optimized a kicking sequence such that we performed




Figure 6.11: Experimentally measured time evolution of the momentum contrast figure of merit, f_p after one and three kicks at $K = 1.06$. The time (horizontal) axis is calibrated in microseconds.




Figure 6.12: Experimentally measured time evolution of the momentum contrast figure of merit, f_p after one and two kicks at $K = 3.2$. The time (horizontal) axis is calibrated in microseconds.

as many kicks as possible such that the time between subsequent kicks was at least $1 \mu\text{s}$. The sequences were optimized empirically, by using the momentum contrast measurements. After each kick, we performed this measurement at a range of times, and placed the next kick at the time of optimal compression as measured by f_p . The results of these measurements are shown in Figs. 6.10, 6.11, and 6.12. For the case of $K = 0.43$, ten kicks occurred before

the kick spacing approached $1 \mu\text{s}$. The spacings between the kicks were $10.0 \mu\text{s}$, $6.9 \mu\text{s}$, $4.7 \mu\text{s}$, $3.9 \mu\text{s}$, $3.0 \mu\text{s}$, $2.2 \mu\text{s}$, $2.1 \mu\text{s}$, $1.8 \mu\text{s}$, and $1.7 \mu\text{s}$. While conducting the experiments it appeared that the overall sequence was fairly sensitive to small fluctuations in the system parameters (e.g., laser power and frequency). These parameters can alter the optimization time after each kick in the sequence, so it is not surprising that a long sequence exhibits this sensitivity. In Fig. 6.10 it is clear that with each subsequent kick, the minimum value of f_p monotonically decreases. For $K = 1.06$, it was only possible to apply three kicks before the compression time approached $1 \mu\text{s}$. The kick spacings for this case were $9.4 \mu\text{s}$ and $3.9 \mu\text{s}$. The third kicking strength was $K = 3.2$. In this case, the longest sequence consisted of two kicks, separated by $4.5 \mu\text{s}$

The values of f_p that we have measured are not spectacular— they are well below those that we expect from simple simulations. Such a simulation, for a single kick with $K = 3.2$, is shown in Fig. 6.7. The minimum of the predicted f_p curve is approximately 0.47. However, the value from our experiment is closer to 0.71. As in many of our other experiments, it is difficult to quantitatively measure many of the values that are extracted from the distributions. We can achieve great precision when measuring at what value of a system parameter the dynamics of the system change. However, we cannot in the cases of measuring the system energy $\langle p^2 \rangle$ or the average momentum $\langle p \rangle$. There are several contributing reasons for this. First, these measurements weight the distribution most highly at high momentum values, where the signal is the noisiest. More importantly, there are nonuniformities in our detection system that tend to decrease the apparent signal when parts of the distribution enter higher momentum regions (see §3.6.2). These effects tend to move the apparent value of $\langle p \rangle$ towards zero in any measurement, moving f_p closer to unity. The magnitude of this effect is doubled because we are looking at the difference between two average momentum values with opposite sign. These effects may well be enough to account for the difference between our measurements and what we might expect from the simple simulations. Using only the locations of the *peaks* of the distributions results in a slightly less biased measurement, but this is only appropriate when the distribution is




Figure 6.13: Pendulum oscillation measurements, $K = 3.2$. We project the distribution into the pendulum potential and allow the distribution to evolve for a total of $27.2 \mu\text{s}$, in 35 steps. At each time step we measure the momentum distribution and compute the Süßmann measure δ_{Sp} , which is plotted here as the vertical (color) variable. Brighter colors correspond to higher values of δ_{Sp} . The time evolution is measured independently at several values of the well depth α_p to trace out the 2D parameter space. This set of measurements is executed with zero, one and two focusing kicks. We are looking for a decrease in the amplitude of the oscillations at some well depth. It is difficult to discern these changes directly from this type of data. These changes are visible in the standard deviation σ_S of each time evolution sequence, as shown in Fig. 6.16.

already well focused. We should note from simulations that the *shape* of the $f(t)$ curve is very sensitive to the exact value of K . It is very interesting that the shape (but not amplitude) of the theoretical and experimental curves shown respectively in Figs. 6.7 and 6.12 matches up. This strengthens the claim that the discrepancy is due to biases in the detection system. Despite the discrepancy in the absolute value, the momentum contrast measurements provide us with a mechanism for observing the gradual increase of the spatial focusing. Furthermore, these measurements provide a practical mechanism for optimizing the time between kicks.

The next major class of data is the pendulum oscillation measurement. The purpose of this data is to provide an independent verification that spatial compression has occurred. As opposed to the f_p data where we try to extract quantitative information with analysis of the distributions, in this case we are looking for a particular value of a parameter (the well depth) where the oscillation amplitude is minimized. We begin by tracing out the two-dimensional parameter space of the evolution time and the well depth. At each point in the




Figure 6.14: Pendulum oscillation measurements, $K = 0.43$. We project the distribution into the pendulum potential at a given well depth and allow the distribution to evolve for $102 \mu\text{s}$ in 35 time steps. At each time step the Süßmann measure δ_{SP} is determined from the momentum distribution. The standard deviation σ_S of the $\delta_{SP}(t)$ is plotted here as a function of the well depth α_p . This procedure is repeated after 0, 2, 4, 6, 8, and 10 kick focusing sequences. In each case, the measurement begins at the time that the next kick *would* occur, if there were to be additional kicks in the sequence. As the number of kicks increases, the degree of spatial focusing increases. This is visible as a reduction in the amplitude of oscillations, characterized by σ_S . For the last case (after $10 \mu\text{s}$), the dip in σ_S is rather broad and spans the entire range of α_p that we are sweeping over.

measurement we compute the Süßmann measure δ_{SP} of the momentum distribution. An example set of data taken with this procedure is shown in Fig. 6.13. From these plots, we can see that there are oscillations in δ_{SP} at every well depth, with a frequency that depends upon the well depth. Visual inspection of these plots is not sufficient to determine where the amplitude of the oscillations is minimized, so we calculate the standard deviation of $\delta_{SP}(t)$ and plot this as a function of the well depth. The results of these measurements are shown in Figs. 6.14, 6.15, and 6.16. For most of the cases with multiple kicks, we can see a clear minimum in the $\sigma_S(\alpha)$ curve.




Figure 6.15: Pendulum oscillation measurements, $K = 1.06$. We project the distribution into the pendulum potential as in Fig. 6.14 and plot σ_S of $\delta_S p(t)$ as a function of the well depth α_p . The procedure is repeated without any focusing kicks, with one kick, and with three kicks. The total evolution time at each point on this plot is $85 \mu\text{s}$. After three kicks, the minimum occurs at $\alpha_p = 3.3$.




Figure 6.16: Pendulum oscillation measurements, $K = 3.2$. We project the distribution into the pendulum potential as in Fig. 6.14 and plot σ_S of $\delta_S p(t)$ as a function of the well depth α_p . The procedure is repeated without any focusing kicks, with one kick, and with two kicks. The total evolution time at each point on this plot is $27.2 \mu\text{s}$. The data shown in this plot is the same as that shown in Fig. 6.13. Note that the vertical scale for the cases with kicks is enlarged relative to the case without any kicks. After one kick, the minimum occurs at $\alpha_p = 8.25$, after two kicks, the minimum occurs at $\alpha_p = 15.5$.

The most interesting cases are those where the focusing is strongest. Let us consider two of these cases, at $K = 1.06$ after three kicks, and at $K = 3.2$ with two kicks. For $K = 1.06$ with three kicks, the minimum of $\sigma_S(\alpha)$ is fitted to occur at $\alpha_p = 3.3$. The width of the corresponding lattice ground state is then $\sigma_x = \sqrt{1/2\sqrt{\alpha}} = 0.53$. In physical units ($x = x'\lambda/4\pi$), this corresponds to a width of about 35 nm. For $K = 3.2$ and two kicks, the minimum occurs at $\alpha_p = 15.5$, corresponding to a width of about 24 nm. This value is comparable to the best feature sizes that have been demonstrated with other atom




Figure 6.17: Pendulum rotation measurements, $K = 1.06$. The two curves show the results of the measurement procedure after 1 kick (blue) and 3 kicks (red). The peaks on the edges of the distribution correspond to the uniform component of the position distribution. The size of this contribution is reduced by additional kicks.

lithography techniques. It is important to recall that our experiments are limited by the response time of the EOM, which is not a fundamental limitation of this technique. Besides the width of the features, we are also very interested to know how much of the distribution is contained in the unfocused background. From the various simulations in this chapter, it is apparent that more than just two or three kicks are required to produce substantial suppression of the background. We can see the size of the background by looking at the amplitude of residual oscillations when $\sigma_S(\alpha)$ is minimized. The lowest value of in our experiments is $\sigma_S = 0.66$, which occurs for $K = 1.06$ with three kicks.

Finally, we come to the results from the free expansion measurements. Let us first consider the results from the simple temperature measurements after one focusing kick. We




Figure 6.18: Pendulum rotation measurements, $K = 0.43$. The two curves show the results of the measurement procedure after 1 kick (green) 5 kicks (blue) and 10 kicks (red). Again, it is apparent that additional kicks serve to reduce the uniform portion of the spatial distribution.

measured the temperature after a single focusing kick in the usual way. For one kick at $K = 0.43$ or $K = 1.06$, the final temperature was about 800 nK or $1.5 \mu\text{K}$, respectively. For one kick at $K = 3.2$, the final temperature was measured to be $7.3 \mu\text{K}$, although this not meaningful because the strongly kicked distribution is not thermal. For a strong enough single kick, the final momentum distribution is dominated by $K\langle\cos x\rangle$ (as in Eq. 6.7) rather than any thermal component. We can also attempt to visualize the position distribution directly by allowing the atoms to evolve for $1/4$ of the harmonic oscillator period of a deep lattice. We have performed this measurement for kick sequences of different lengths for the two weaker kick strengths. (For the case with $K = 3.2$, the kick sequence is very short and we did not perform the rotation measurement.) The distributions that are produced by this method, shown in Figs. 6.17 and 6.18, are suggestive of those that appear in Fig. 6.5. There

is a slight asymmetry in the momentum distributions that is probably due to nonuniformities in the detection system. It is also possible that a small but nonzero initial velocity of the atomic distribution could lead to contribute to this effect.

We have fitted the data in Figs. 6.17 and 6.18 to determine the relative amplitudes of the central Gaussian portion and the side-peak portion that corresponds to the uniform component of the atomic distribution. For consistency of the fitting procedure, the momentum distributions were first symmetrized by averaging each distribution with its symmetric reflection about $p = 0$. A Gaussian curve was fit locally to the distributions near the $p = 0$ and separately near $p = \pm 18$, where the uniform component of the distribution is peaked. From these two fits, we determine the peak height of these two parts of the curve. For the case of $K = 1.06$, after one kick, the peaks from the uniform component have an amplitude that is 59% of the central peak. After two kicks, the side peaks have an amplitude that is 30% of the central peak. For the kick sequence with $K = 0.43$, the relative uniform peak height was determined to be 83%, 57%, and 45% after 1, 5, and 10 kicks. In both cases, it is clear that the relative size of the uniform component is decreasing. It may be possible to combine the results of this measurement with the results of the pendulum oscillation measurements to determine the absolute magnitude of the residual background, which we hope to accomplish for a future publication. Part of the difficulty with this analysis is that the height of the peaks depends sensitively upon the width of the distribution, as we saw in Figs. 6.8 and 6.9.

In summary, we have observed the spatial compression of an atomic distribution by a new technique. The degree of focusing is comparable to the best that has been observed in other laboratories by focusing atoms with a single standing wave of light. In our case, the repeated application of pulses leads to increased focusing of the distribution, at the expense of heating. But beyond this simple focusing, we have preliminary indications of suppression of the uniform background that has been a problem for other atom lithography techniques. To achieve the best possible background suppression requires a sequence of kicks that continues until the kicks are nearly overlapping. The EOM response time in our

system is a more fundamental limitation and we do not expect to be able to observe a high degree of background suppression in our setup. However, this is only a limitation of our detection system and should not be a limitation for applications of this technology. It should be straightforward to implement this focusing method in an experimental setup that can directly deposit the atoms onto a substrate.

Bibliography

- [ALLEN87] L. C. Allen and J. H. Eberly, *Optical Resonance and Two-Level Atoms* (Dover, New York, 1987).
- [ALTLAND96] Alexander Altland and Martin R. Zirnbauer, “Field Theory of the Quantum Kicked Rotor,” *Phys. Rev. Lett.* **77**, 4536 (1996).
- [AMMANN98] H. Ammann, R. Gray, I. Shvarchuck, and N. Christensen, “Quantum Delta-Kicked Rotor: Experimental Observation of Decoherence,” *Phys. Rev. Lett.* **80**, 4111 (1998).
- [ANDERSON58] P. W. Anderson, “Absence of Diffusion in Certain Random Lattices,” *Phys. Rev.* **109**, 1492 (1958).
- [ANDERSON99] W. R. Anderson, C. C. Bradley, J. J. McClelland, and R. J. Celotta, “Minimizing feature width in atom optically fabricated chromium nanostructures,” *Phys. Rev. A* **59**, 2476 (1999).
- [ARNDT91] M. Arndt, A. Buchleitner, R. N. Mantegna, and H. Walther, “Experimental Study of Quantum and Classical Limits in Microwave Ionization of Rubidium Rydberg Atoms,” *Phys. Rev. Lett.* **67**, 2435 (1991).
- [ARNDT97] M. Arndt, M. Ben Dahan, D. Guéry-Odelin, M. W. Reynolds, and J. Dalibard, “Observation of a Zero-Energy Resonance in Cs-Cs Collisions,” *Phys. Rev. Lett.* **79**, 625 (1997).
- [AVERBUKH01] I. Sh. Averbukh and R. Arvieu, “Angular Focusing, Squeezing and Rainbow Formation in a Strongly Driven Quantum Rotor,” *Phys. Rev. Lett.* **87**, 163601 (2001).
- [BALYKIN88] V. I. Balykin, V. S. Letokhov, Yu. B. Ovchinnikov, and A. I. Sidorov, “Quantum-State-Selective Mirror Reflection of Atoms by Laser Light,” *Phys. Rev. Lett.* **60**, 2137 (1988).
- [BAYFIELD74] J. E. Bayfield and P. M. Koch, “Multiphoton Ionization of Highly Excited Hydrogen Atoms,” *Phys. Rev. Lett.* **33**, 258 (1974).

- [BAYFIELD91] James E. Bayfield, “Near-classical noise enhancement of microwave ionization of Rydberg atoms,” *Chaos* **1**, 110 (1991).
- [BERGGREN95] Karl K. Berggren, Andreas Bard, James L. Wilbur, John D. Gillaspay, Andreas G. Helg, Jabez J. McClelland, Steven L. Rolston, William D. Phillips, Mara Prentiss, and George W. Whitesides, “Laser-Focused Atomic Deposition,” *Science* **269**, 1255 (1995).
- [BERRY99] M. V. Berry and E. Bodenschatz, “Caustics, multiply reconstructed by Talbot interference,” *J. Mod. Opt.* **46**, 349 (1999).
- [BHARUCHA97] Cyrus Farrokh Bharucha, *Experiments in Dynamical Localization of Ultra-Cold Sodium Atoms Using Time-Dependent Optical Potentials*, Ph.D. dissertation, The University of Texas at Austin (1997).
- [BHARUCHA99] C. F. Bharucha, J. C. Robinson, F. L. Moore, Qian Niu, Bala Sundaram, and M. G. Raizen, “Dynamical localization of ultracold sodium atoms,” *Phys. Rev. E* **60**, 3881 (1999).
- [BJORKHOLM78] J. Bjorkholm, R. Freeman, A. Ashikin, and D. Pearson, “Observation of Focussing of Neutral Atoms by the Dipole Forces of Resonance-Radiation Pressure,” *Phys. Rev. Lett.* **41**, 1361 (1978).
- [BJORKLUND83] G. C. Bjorklund, M. D. Levenson, W. Lenth, and C. Ortiz, “Frequency Modulation (FM) Spectroscopy,” *Appl. Phys. B* **32**, 145 (1983).
- [BOHIGAS93] O. Bohigas, S. Tomsovic, and D. Ullmo, “Manifestations of Classical Phase Space Structures in Quantum Mechanics,” *Phys. Rep.* **223**, 43 (1993).
- [BOIRON96] D. Boiron, A. Michaud, P. Lemonde, Y. Castin, C. Salomon, S. Weyers, K. Szymaniec, L. Cagnet, and A. Clairon, “Laser cooling of cesium atoms in gray optical molasses down to 1.1 μK ,” *Phys. Rev. A* **53**, R3734 (1996).
- [BRUN96] Todd A. Brun, Ian C. Percival, and Rüdiger Schack, “Quantum chaos in open systems: a quantum state diffusion analysis,” *J. Phys. A* **29**, 2077 (1996).
- [CASATI79] G. Casati, B. V. Chirikov, F. M. Izraelev, and Joseph Ford, “Stochastic Behavior of a Quantum Pendulum under a Periodic Perturbation,” in *Stochastic Behavior in Classical and Quantum Hamiltonian Systems: Proceedings of the Volta Memorial Conference, Como, 1977 (Lecture Notes in Physics Vol. 93)*, G. Casati and J. Ford, Eds. (Springer-Verlag, Berlin, 1979).

- [CASATI94] Giulio Casati, Robert Graham, Italo Guarneri, and Felix M. Izrailev, “Tunneling between localized states in classically chaotic systems,” *Phys. Lett. A* **190**, 159 (1994).
- [CASTIN98] Y. Castin, J. I. Cirac, and M. Lewenstein, “Reabsorption of Light by Trapped Atoms,” *Phys. Rev. Lett.* **80**, 5305 (1998).
- [CHAPMAN95] Michael S. Chapman, Troy D. Hammond, Alan Lenef, Jörg Schmiedmayer, Richard A. Rubenstein, Edward Smith, and David E. Pritchard, “Photon scattering from atoms in an atom interferometer: Coherence lost and regained,” *Phys. Rev. Lett.* **75**, 3783 (1995).
- [CHEN92] Jian Chen, J. G. Story, J. J. Tollett, and Randall G. Hulet, “Adiabatic cooling of atoms by an intense standing wave,” *Phys. Rev. Lett.* **69**, 1344 (1992).
- [CHIRIKOV79] Boris V. Chirikov, “A Universal Instability of Many-Dimensional Oscillator Systems,” *Phys. Rep.* **52**, 263 (1979).
- [CHIRIKOV84] B. V. Chirikov and D. L. Shepelyansky, “Correlation Properties of Dynamical Chaos in Hamiltonian Systems,” *Physica D* **13**, 395 (1984).
- [CHIRIKOV95] B. V. Chirikov and D. L. Shepelyansky, “Shnirelman Peak in Level Spacing Statistics,” *Phys. Rev. Lett.* **74**, 518 (1995).
- [CHU96] A. P. Chu, K. K. Berggren, K. S. Johnson, and M. G. Prentiss, “A virtual slit for atom optics and nanolithography,” *Quant. Semiclass. Opt.* **8**, 521 (1996).
- [CIRAC96] J. I. Cirac, M. Lewenstein, and P. Zoller, “Collective laser cooling of trapped atoms,” *Europhys. Lett.* **35**, 647 (1996).
- [CLARKE95] R. M. Clarke, I. H. Chan, C. M. Marcus, C. I. Duruöz, J. S. Harris, Jr., K. Campman, and A. C. Gossard, “Temperature dependence of phase breaking in ballistic quantum dots,” *Phys. Rev. B* **52**, 2656 (1995).
- [COHEN91] Doron Cohen, “Quantum chaos, dynamical correlations, and the effect of noise on localization,” *Phys. Rev. A* **44**, 2249 (1991).
- [COHEN99] Doron Cohen, “Non Perturbative Destruction of Localization in the Quantum Kicked Particle Problem,” arXiv.org preprint [chaos-dyn/9909016](https://arxiv.org/abs/chaos-dyn/9909016) (1999).
- [DALIBARD89] J. Dalibard and C. Cohen-Tannoudji, “Laser cooling below the Doppler limit by polarization gradients: simple theoretical models,” *J. Opt. Soc. Am. B* **6**, 2023 (1989).

- [DAVIS81] M. J. Davis and E. J. Heller, “Quantum Dynamical Tunneling in Bound States,” *J. Chem. Phys.* **75**, 246 (1981).
- [DELANDE01] D. Delande, “Quantum Chaos in Atomic Physics,” in *Coherent Atomic Matter Waves: Proceedings of the Les Houches Summer School, Session LXXII, 27 July – 27 August 1999*, R. Kaiser, C. Westbrook, and F. David, Eds. (Springer-Verlag, Berlin, 2001).
- [DEMBOWSKI00] C. Dembowski, H.-D. Gräf, A. Heine, R. Hofferbert, H. Rehfeld, and A. Richter, “First Experimental Evidence for Chaos-Assisted Tunneling in a Microwave Annular Billiard,” *Phys. Rev. Lett.* **84**, 867 (2000).
- [DEMTRÖDER96] W. Demtröder, *Laser Spectroscopy*, second ed. (Springer-Verlag, Berlin, 1996).
- [DEPUE99] Marshall T. DePue, Colin McCormick, S. Lukman Winoto, Steven Oliver, and David S. Weiss, “Unity Occupation of Sites in a 3D Optical Lattice,” *Phys. Rev. Lett.* **82**, 2262 (1999).
- [DIDDAMS00] Scott A. Diddams, David J. Jones, Jun Ye, Steven T. Cunliff, John L. Hall, Jinendra K. Ranka, Robert S. Windeler, Ronald Holzwarth, Thomas Udem, and T. W. Hänsch, “Direct Link between Microwave and Optical Frequencies with a 300 THz Femtosecond Laser Comb,” *Phys. Rev. Lett.* **84**, 5102 (2000).
- [FISCHER93] Martin Christian Fischer, *Design and Performance of a Ring Dye Laser*, Master’s thesis, The University of Texas at Austin (1993).
- [FISCHER00A] Baruch Fischer, Amir Rosen, Alexander Bekker, and Shmuel Fishman, “Experimental observation of localization in the spatial frequency domain of a kicked optical system,” *Phys. Rev. E* **61**, R4694 (2000).
- [FISCHER00B] Baruch Fischer, Amir Rosen, Alexander Bekker, and Shmuel Fishman, “Experimental observation of localization in the spatial frequency domain of a kicked optical system,” *Phys. Rev. E* **61**, R4694 (2000).
- [FISHMAN82] Shmuel Fishman, D. R. Grempel, and R. E. Prange, “Chaos, Quantum Recurrences, and Anderson Localization,” *Phys. Rev. Lett.* **49**, 509 (1982).
- [FRIEDBURG51] H. Friedburg and W. Paul, “Optical Presentation with Neutral Atoms,” *Naturwissenschaften* **38**, 159 (1951).
- [FRIEDRICH95] Bretislav Friedrich and Dudley Herschbach, “Alignment and Trapping of Molecules in Intense Laser Fields,” *Phys. Rev. Lett.* **74**, 4623 (1995).

- [GIBBLE93] Kurt Gibble and Steven Chu, “Laser-cooled Cs frequency standard and a measurement of the frequency shift due to ultracold collisions,” *Phys. Rev. Lett.* **70**, 1771 (1993).
- [GILTNER95] David M. Giltner, Roger W. McGowan, and Siu Au Lee, “Theoretical and experimental study of the Bragg scattering of atoms from a standing light wave,” *Phys. Rev. A* **52**, 3966 (1995).
- [GODUN00] R. M. Godun, M. B. d’Arcy, M. K. Oberthaler, G. S. Summy, and K. Burnett, “Quantum accelerator modes: A tool for atom optics,” *Phys. Rev. A* **62**, 013411 (2000).
- [GORDON80] J. P. Gordon and A. Ashkin, “Motion of atoms in a radiation trap,” *Phys. Rev. A* **21**, 1606 (1980).
- [GRAHAM92] R. Graham, M. Schlautmann, and P. Zoller, “Dynamical localization of atomic-beam deflection by a modulated standing light wave,” *Phys. Rev. A* **45**, R19 (1992).
- [GROSSMANN93] Frank Grossmann, Thomas Dittrich, Peter Jung, and Peter Hänggi, “Coherent Transport in a Periodically Driven Bistable System,” *J. Stat. Phys.* **70**, 229 (1993).
- [GUPTA95] R. Gupta, J. J. McClelland, Z. J. Jabbour, and R. J. Celotta, “Nanofabrication of a two-dimensional array using laser-focused atomic deposition,” *Appl. Phys. Lett.* **67**, 1378 (1995).
- [GUPTA96] R. Gupta, J. J. McClelland, P. Marte, and R. J. Celotta, “Raman-Induced Avoided Crossings in Adiabatic Optical Potentials: Observation of $\lambda/8$ Spatial Frequency in the Distribution of Atoms,” *Phys. Rev. Lett.* **76**, 4689 (1996).
- [HABIB98] Salman Habib, Hideo Mabuchi, Kozuke Shizume, and Bala Sundaram, “Comment on ‘Quantum Delta-Kicked Rotor: Experimental Observation of Decoherence’,” unpublished (1998).
- [HALL76] J. L. Hall and S. A. Lee, “Interferometric real-time display of CW dye laser wavelength with sub-Doppler accuracy,” *Appl. Phys. Lett.* **29**, 367 (1976).
- [HAMANN98] S. E. Hamann, D. L. Haycock, G. Klose, P. H. Pax, I. H. Deutsch, and P. S. Jessen, “Resolved-Sideband Raman Cooling to the Ground State of an Optical Lattice,” *Phys. Rev. Lett.* **80**, 4149 (1998).
- [HÄNSCH80] T. W. Hänsch and B. Couillaud, “Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity,” *Opt. Comm.* **35**, 441 (1980).

- [HANSON84] J. D. Hanson, E. Ott, and T. M. Antonsen, “Influence of finite wavelength on the quantum kicked rotator in the semiclassical regime,” *Phys. Rev. A* **29**, 819 (1984).
- [HEMMERICH94] A. Hemmerich, M. Weidemler, T. Esslinger, C. Zimmermann, and T. Hänsch, “Trapping Atoms in a Dark Optical Lattice,” *Phys. Rev. Lett.* **53**, 37 (1994).
- [HENSINGER01] W. K. Hensinger, H. Häffner, A. Browaeys, N. R. Heckenberg, K. Helmerston, C. McKenzie, G. J. Milburn, W. D. Phillips, S. L. Rolston, H. Rubinsztein-Dunlop, and B. Upcroft, “Dynamical tunneling of ultracold atoms,” *Nature* **412**, 52 (2001).
- [HUG01] M. Hug and G. J. Milburn, “Quantum slow motion,” *Phys. Rev. A* **63**, 023413 (2001).
- [IZRAILEV79] F. M. Izrailev and D. L. Shepelyanskii, “Quantum resonance for a rotor in a nonlinear periodic field,” *Sov. Phys. Dokl.* **24**, 996 (1979).
- [KARNEY83] Charles F. F. Karney, “Long-Time Correlations in the Stochastic Regime,” *Physica D* **8**, 585 (1983).
- [KASEVICH91] Mark Kasevich and Steven Chu, “Atomic interferometry using stimulated Raman transitions,” *Phys. Rev. Lett.* **67**, 181 (1991).
- [KASEVICH92] Mark Kasevich and Steven Chu, “Laser cooling below a photon recoil with three-level atoms,” *Phys. Rev. Lett.* **69**, 1741 (1992).
- [KASTBERG95] A. Kastberg, W. D. Phillips, S. L. Rolston, and R. J. C. Spreeuw, “Adiabatic Cooling of Cesium to 700 nK in an Optical Lattice,” *Phys. Rev. Lett.* **74**, 1542 (1995).
- [KEITH91] David W. Keith, Christopher R. Ekstrom, Quentin A. Turchette, and David E. Pritchard, “An interferometer for atoms,” *Phys. Rev. Lett.* **66**, 2693 (1991).
- [KETTERLE93] Wolfgang Ketterle, Kendall B. Davis, Michael A. Joffe, Alex Martin, and David E. Pritchard, “High Densities of Cold Atoms in a *Dark* Spontaneous-Force Optical Trap,” *Phys. Rev. Lett.* **70**, 2253 (1993).
- [KHAYKOVICH00] L. Khaykovich and N. Davidson, “Adiabatic focusing of cold atoms in a blue-detuned laser standing wave,” *Appl. Phys. B* **70**, 683 (2000).
- [KIELPINSKI01] D. Kielpinski, V. Meyer, M. A. Rowe, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland, “A Decoherence-Free Quantum Memory Using Trapped Ions,” *Science* **291**, 1013–1015 (2001).
- [KING99] Brian E. King, *Quantum State Engineering and Information Processing with Trapped Ions*, Ph.D. thesis, University of Colorado at Boulder (1999).

- [KLAFTER96] Joseph Klafter, Michael F. Schlesinger, and Gert Zumofen, “Beyond Brownian Motion,” *Physics Today* p. 33 (February 1996).
- [KLAPPAUF98A] B. G. Klappauf, W. H. Oskay, D. A. Steck, and M. G. Raizen, “Experimental Study of Quantum Dynamics in a Regime of Classical Anomalous Diffusion,” *Phys. Rev. Lett.* **81**, 4044 (1998).
- [KLAPPAUF98B] B. G. Klappauf, W. H. Oskay, D. A. Steck, and M. G. Raizen, “Observation of Noise and Dissipation Effects on Dynamical Localization,” *Phys. Rev. Lett.* **81**, 1203 (1998). Also **82**, 241(E) (1998).
- [KLAPPAUF98C] Bruce George Klappauf, *Experimental Studies of Quantum Chaos with Trapped Cesium*, Ph.D. dissertation, The University of Texas at Austin (1998).
- [KLAPPAUF99] B. G. Klappauf, W. H. Oskay, D. A. Steck, and M. G. Raizen, “Quantum Chaos with Cesium Atoms: Pushing the Boundaries,” *Physica D* **131**, 78 (1999).
- [KOCH95] Peter M. Koch, “Microwave ‘ionization’ of excited hydrogen atoms: How nonclassical local stability brought about by scarred separatrix states is affected by broadband noise and by varying the pulse envelope,” *Physica D* **83**, 178 (1995).
- [KOKOROWSKI01] David A. Kokorowski, Alexander D. Cronin, Tony D. Roberts, and David E. Pritchard, “From Single- to Multiple-Photon Decoherence in an Atom Interferometer,” *Phys. Rev. Lett.* **86**, 2191 (2001).
- [KOZUMA99] M. Kozuma, L. Deng, E. W. Hagley, J. Wen, R. Lutwak, K. Helmerson, S. L. Rolston, and W. D. Phillips, “Coherent Splitting of Bose-Einstein Condensed Atoms with Optically Induced Bragg Diffraction,” *Phys. Rev. Lett.* **82**, 871 (1999).
- [LICHTENBERG92] A. J. Lichtenberg and M. A. Leiberman, *Regular and Chaotic Dynamics*, second ed. (Springer-Verlag, New York, 1992).
- [LIN90] W. A. Lin and L. E. Ballentine, “Quantum Tunneling and Chaos in a Driven Anharmonic Oscillator,” *Phys. Rev. Lett.* **65**, 2927 (1990).
- [MACADAM92] K. B. MacAdam, A. Steinbach, and C. Wieman, “A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb,” *Am. J. Phys.* **60**, 1098 (1992).
- [MADISON98] Kirk William Madison, *Quantum Transport in Optical Lattices*, Ph.D. dissertation, The University of Texas at Austin (1998).
- [MAIN91] J. Main, G. Wiebusch, and K. H. Welge, “Spectroscopy of the classically chaotic hydrogen atom in magnetic fields,” *Comm. At. Mol. Phys.* **25**, 233 (1991).

- [MARCUS92] C. M. Marcus, A. J. Rimberg, R. M. Westervelt, P. F. Hopkins, and A. C. Gosard, "Conductance Fluctuations and Chaotic Scattering in Ballistic Microstructures," *Phys. Rev. Lett.* **69**, 506 (1992).
- [MARTIN87] P. J. Martin, P. L. Gould, B. G. Oldaker, A. H. Miklich, and D. E. Pritchard, "Diffraction of atoms moving through a standing light wave," *Phys. Rev. A* **36**, 2495 (1987).
- [MARTIN88] Peter J. Martin, Bruce G. Oldaker, Andrew H. Miklich, and David E. Pritchard, "Bragg scattering of atoms from a standing light wave," *Phys. Rev. Lett.* **60**, 515 (1988).
- [MCCLELLAND93] J. J. McClelland, R. E. Scholten, E. C. Palm, and R. J. Celotta, "Laser-Focused Atomic Deposition," *Science* **262**, 877 (1993).
- [MCCLELLAND98] J. J. McClelland, R. Gupta, R. J. Celotta, and G. A. Porkolab, "Nanostucture fabrication by reactive-ion etching of laser-focused chromium on silicon," *Appl. Phys. B* **66**, 95 (1998).
- [MCGOWAN95] Roger W. McGowan, David M. Giltner, and Siu Au Lee, "Light force cooling, focusing, and nanometer-scale deposition of aluminum atoms," *Opt. Lett.* **20**, 2535 (1995).
- [METCALF99] Harold J. Metcalf and Peter van der Straten, *Laser cooling and trapping* (Springer-Verlag, New York, 1999).
- [MIGDALL85] Alan L. Migdall, John V. Prodan, William D. Phillips, Thomas H. Bergeman, and Harold J. Metcalf, "First Observation of Magnetically Trapped Neutral Atoms," *Phys. Rev. Lett.* **54**, 2596 (1985).
- [MILNER00] V. Milner, D. A. Steck, W. H. Oskay, and M. G. Raizen, "Recovery of Classically Chaotic Behavior in a Noise-Driven Quantum System," *Phys. Rev. E* **61**, 7223 (2000).
- [MOLER92] Kathryn Moler, David S. Weiss, Mark Kasevich, and Steven Chu, "Theoretical analysis of velocity-selective Raman transitions," *Phys. Rev. A* **45**, 342 (1992).
- [MONROE90] C. Monroe, W. Swann, H. Robinson, and C. Wieman, "Very cold trapped atoms in a vapor cell," *Phys. Rev. Lett.* **65**, 1571 (1990).
- [MONROE92] Christopher Roy Monroe, *Experiments with Optically and Magnetically Trapped Cesium Atoms*, Ph.D. thesis, University of Colorado (1992).

- [MONROE96] C. Monroe, D. M. Meekhof, B. E. King, and D. J. Wineland, “A ‘Schrödinger Cat’ Superposition State of an Atom,” *Science* **272**, 1131 (1996).
- [MOORE94] F. L. Moore, J. C. Robinson, C. Bharucha, P. E. Williams, and M. G. Raizen, “Observation of Dynamical Localization in Atomic Momentum Transfer: A New Testing Ground for Quantum Chaos,” *Phys. Rev. Lett.* **73**, 2794 (1994).
- [MOORE95] F. L. Moore, J. C. Robinson, C. F. Bharucha, Bala Sundaram, and M. G. Raizen, “Atom Optics Realization of the Quantum δ -Kicked Rotor,” *Phys. Rev. Lett.* **75**, 4598 (1995).
- [MOSKOWITZ83] Philip E. Moskowitz, Phillip L. Gould, Susan R. Atlas, and David E. Pritchard, “Diffraction of an Atomic Beam by Standing-Wave Radiation,” *Phys. Rev. Lett.* **51**, 370 (1983).
- [MOUCHET01] A. Mouchet, C. Miniatura, R. Kaiser, B. Grémaud, and D. Delande, “Chaos assisted tunnelling with cold atoms,” *Phys. Rev. E* **64**, 016221 (2001).
- [MÜTZEL00] M. Mützel, D. Haubrich, and D. Meschede, “Nanoscale focusing of atoms with a pulsed standing wave,” *Appl. Phys. B* **70**, 689 (2000).
- [NESMEYANOV63] A. N. Nesmeyanov, *Vapor Pressure of the Chemical Elements* (Elsevier, Amsterdam, 1963). English edition edited by Robert Gary.
- [NOWAK97] S. Nowak, Ch. Kurtsiefer, and T. Pfau, “High-order Talbot fringes for atomic matter waves,” *Opt. Lett.* **22**, 1430 (1997).
- [OBERTHALER99] M. K. Oberthaler, R. M. Godun, M. B. d’Arcy, G. S. Summy, and K. Burnett, “Observation of Quantum Accelerator Modes,” *Phys. Rev. Lett.* **83**, 4447 (1999).
- [OSKAY00] W. H. Oskay, D. A. Steck, V. Milner, B. G. Klappauf, and M. G. Raizen, “Ballistic Peaks at Quantum Resonance,” *Opt. Comm.* **179**, 137 (2000).
- [OTT84] E. Ott, T. M. Antonsen, Jr., and J. D. Hanson, “Effect of Noise on Time-Dependent Quantum Chaos,” *Phys. Rev. Lett.* **53**, 2187 (1984).
- [PAPPAS80] P. G. Pappas, M. M. Burns, D. D. Hinshelwood, M. S. Feld, and D. E. Murnick, “Saturation spectroscopy with laser optical pumping in atomic barium,” *Phys. Rev. A* **21**, 1955 (1980).
- [PAZ01] J. P. Paz and W. H. Zurek, “Decoherence,” in *Coherent Atomic Matter Waves: Proceedings of the Les Houches Summer School, Session LXXII, 27 July – 27 August 1999*, R. Kaiser, C. Westbrook, and F. David, Eds. (Springer-Verlag, Berlin, 2001).

- [PHILLIPS82] William D. Phillips and Harold Metcalf, "Laser Deceleration of an Atomic Beam," *Phys. Rev. Lett.* **48**, 596 (1982).
- [RAAB87] E. L. Raab, M. Prentiss, Alex Cable, Steven Chu, and D. E. Pritchard, "Trapping of neutral sodium atoms with radiation pressure," *Phys. Rev. Lett.* **59**, 2631 (1987).
- [RECHESTER80] A. B. Rechester and R. B. White, "Calculation of Turbulent Diffusion for the Chirikov-Taylor Model," *Phys. Rev. Lett.* **44**, 1586 (1980).
- [RECHESTER81] A. B. Rechester, M. N. Rosenbluth, and R. B. White, "Fourier-space paths applied to the calculation of diffusion for the Chirikov-Taylor model," *Phys. Rev. A* **23**, 2664 (1981).
- [REICHL92] L. E. Reichl, *The Transition to Chaos in Conservative Classical Systems: Quantum Manifestations* (Springer-Verlag, New York, 1992).
- [ROBINSON95A] J. C. Robinson, C. Bharucha, F. L. Moore, R. Jahnke, G. A. Georgakis, Q. Niu, M. G. Raizen, and Bala Sundaram, "Study of Quantum Dynamics in the Transition from Classical Stability to Chaos," *Phys. Rev. Lett.* **74**, 3963 (1995).
- [ROBINSON95B] John Charles Robinson, *Atom Optics: A New Testing Ground for Quantum Chaos*, Ph.D. dissertation, The University of Texas at Austin (1995).
- [ROBINSON96] J. C. Robinson, C. F. Bharucha, K. W. Madison, F. L. Moore, Bala Sundaram, S. R. Wilkinson, and M. G. Raizen, "Can a Single-Pulse Standing Wave Induce Chaos in Atomic Motion?" *Phys. Rev. Lett.* **76**, 3304 (1996).
- [SALOMON90] C. Salomon, J. Dalibard, W. D. Phillips, A. Clairon, and S. Guellati, "Laser Cooling of Cesium Atoms below $3 \mu\text{K}$," *Europhys. Lett.* **12**, 683 (1990).
- [SCHLEICH01] Wolfgang P. Schleich, *Quantum Optics in Phase Space*, first ed. (Wiley-VCH, Berlin, 2001).
- [SESKO91] D. W. Sesko, T. G. Walker, and C. E. Wieman, "Behavior of neutral atoms in a spontaneous force trap," *J. Opt. Soc. Am. B* **8**, 946 (1991).
- [SHEPELYANSKY83] D. L. Shepelyansky, "Some Statistical Properties of Simple Classically Stochastic Quantum Systems," *Physica D* **8**, 208 (1983).
- [SHEPELYANSKY87] D. L. Shepelyansky, "Localization of Diffusive Excitation in Multi-Level Systems," *Physica D* **28**, 103 (1987).
- [SHIOKAWA95] K. Shiokawa and B. L. Hu, "Decoherence, delocalization, and irreversibility in quantum chaotic systems," *Phys. Rev. E* **52**, 2497 (1995).

- [SIRKO93] L. Sirko, M. R. W. Bellermand, A. Haffmans, P. M. Koch, and D. Richards, “Probing Quantal Dynamics of Mixed Phase Space Systems with Noise,” *Phys. Rev. Lett.* **71**, 2895 (1993).
- [SIRKO00] L. Sirko, Sz. Bauch, Y. Hlushchuk, P. M. Koch, R. Blümel, M. Barth, U. Kuhl, and H.-J. Stöckmann, “Observation of dynamical localization in a rough microwave cavity,” *Phys. Lett. A* **266**, 331 (2000).
- [STECK00] Daniel A. Steck, Valery Milner, Windell H. Oskay, and Mark G. Raizen, “Quantitative Study of Amplitude Noise Effects on Dynamical Localization,” *Phys. Rev. E* **62**, 3461 (2000).
- [STECK01A] Daniel A. Steck, “Cesium D Line Data,” available online at <http://www.ph.utexas.edu/~quantopt> (2001).
- [STECK01B] Daniel A. Steck, *Quantum Chaos, Transport, and Decoherence in Atom Optics*, Ph.D. dissertation, The University of Texas at Austin (2001).
- [STECK01C] Daniel A. Steck, Windell H. Oskay, and Mark G. Raizen, “Observation of Chaos-Assisted Tunneling between Islands of Stability,” *Science* (2001).
- [SUNDARAM99] Bala Sundaram and G. M. Zaslavsky, “Anomalous transport and quantum-classical correspondence,” *Phys. Rev. E* **59**, 7231 (1999).
- [SÜSSMANN97] G. Süßmann, “Uncertainty Relation: From Inequality to Equality,” *Z. Naturforsch* **52a**, 49 (1997).
- [TANNER88] Carol E. Tanner and Carl Wieman, “Precision measurement of the hyperfine structure of the ^{133}Cs $6P_{3/2}$ state,” *Phys. Rev. A* **38**, 1616 (1988).
- [TAYLOR95] B. N. Taylor, *Guide for the Use of the International System of Units (SI)*, NIST Special Publication 811, 1995 ed. (US Government Printing Office, Washington DC, 1995).
- [TIMP92] G. Timp, R. E. Behringer, D. M. Tennant, J. E. Cunningham, M. Prentiss, and K. K. Berggren, “Using Light as a Lens for Submicron, Neutral-Atom Lithography,” *Phys. Rev. Lett.* **69**, 1636 (1992).
- [TOMSOVIC94] Steven Tomsovic and Denis Ullmo, “Chaos-assisted tunneling,” *Phys. Rev. E* **50**, 145 (1994).
- [TOMSOVIC98] Steven Tomsovic, “Chaos-assisted tunnelling in the absence of reflexion symmetry,” *J. Phys. A: Math. Gen.* **31**, 9469 (1998).

- [TULCHINSKY98] D. A. Tulchinsky, M. H. Kelley, J. J. McClelland, R. Gupta, and R. J. Celotta, "Fabrication and Domain Imaging of Iron Magnetic Nanowire Arrays," *J. Vac. Sci. Techn. A* **16**, 1817 (1998).
- [VULETIĆ98] Vladan Vuletić, Cheng Chin, Andrew J. Kerman, and Steven Chu, "Degenerate Raman Sideband Cooling of Trapped Cesium Atoms at Very High Atomic Densities," *Phys. Rev. Lett.* **81**, 5768 (1998).
- [WILKINSON96] P. B. Wilkinson, T. M. Fromhold, L. Eaves, F. W. Sheard, N. Miura, and T. Takamasu, "Observation of 'scarred' wavefunctions in a quantum well with chaotic electrons dynamics," *Nature* **380**, 608 (1996).
- [WINOTO99] S. Lukman Winoto, Marshall T. DePue, Nathan E. Bramall, and David S. Weiss, "Laser cooling at high density in deep far-detuned optical lattices," *Phys. Rev. A* **59**, R19 (1999).
- [WOLF00] Steffen Wolf, Steven J. Oliver, and David S. Weiss, "Suppression of Recoil Heating by an Optical Lattice," *Phys. Rev. Lett.* **85**, 4249 (2000).
- [ZASLAVSKY97] G. M. Zaslavsky, M. Edelman, and B. A. Niyazov, "Self-similarity, renormalization, and phase space nonuniformity of Hamiltonian chaotic dynamics," *Chaos* **7**, 159 (1997).
- [ZHONG01] Jianxin Zhong, R. B. Diener, Daniel A. Steck, Windell H. Oskay, Mark G. Raizen, E. Ward Plummer, Zhenyu Zhang, and Qian Niu, "Shape of the Quantum Diffusion Front," *Phys. Rev. Lett.* **86**, 2485 (2001).
- [ZUREK85] Wojciech H. Zurek, "Decoherence and the Transition from Quantum to Classical," *Phys. Today* p. 36 (October 1985). See also the follow-up discussion in *Phys. Today* (April 1993).
- [ZUREK94] Wojciech Hubert Zurek and Juan Pablo Paz, "Decoherence, Chaos, and the Second Law," *Phys. Rev. Lett.* **72**, 2508 (1994).

Vita

Windell Haven Oskay was born on October 26, 1974 in Portland, Oregon, to parents now known as Billy Oskay and Jay Sheckley. He survived Lincoln High School in Portland and graduated in May of 1992. After graduation, he enrolled at Lake Forest College (Lake Forest, IL) and received the degree of Bachelor of Arts in Physics and Mathematics in 1996. In the fall of 1996, he entered the Graduate School of The University of Texas at Austin, where he performed the research described in this dissertation.

Permanent address: 34805 SE Hurlburt Rd.
Corbett, OR 97019

This dissertation was typeset with L^AT_EX[†] by the author.

[†]L^AT_EX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's T_EX Program.