DETECTING SINGLE, TRAPPED IONS

by

Robert Mark Weisskoff

A.B. Physics, Harvard University
(1983)

Submitted to the Department of Physics in
partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May, 1988

Signature of Author _____
Department of Physics
May 6, 1988

Certified by _____
David E. Pritchard
Professor of Physics
Thesis Supervisor

Accepted by _____
George F. Koster
Chairman, Department Committee
on Graduate Studies

Detecting Single, Trapped Ions

by

Robert Mark Weisskoff

Submitted to the Department of Physics
on May 6, 1988, in partial fulfillment of the
requirements for the degree of Doctor of Philosophy

Abstract

We have designed and constructed an experiment to trap a single ion in a Penning trap and measure its motion with an RF SQUID-based, superconducting detector. Using this apparatus we have detected the axial motion of a single, trapped N_2^+ ion.

We describe the theory of ion motion in an imperfect Penning trap and use this theory to explain our measurements on trapped ions. We describe and demonstrate several schemes to detect the motions of trapped ions, including CW and pulsed techniques with one and more drives. We also discuss techniques to cool and manipulate trapped ions (including contaminant ions) and to measure important perturbations from ideal behavior. In addition, we describe several novel measurements of the cyclotron frequency, including a "classical avoided-crossing" and a two-frequency resonance technique. We demonstrate these techniques on small (<10) clouds of N_2^+ ions.

The ultimate goal of the experiment is to measure cyclotron resonances to compare the masses of individual ions at an accuracy of 10^{-11}.

Thesis Supervisor: Dr. David E. Pritchard
Title: Professor of Physics

To Ann Marie

TABLE OF CONTENTS

CHAPTER I:

INTRODUCTION

This thesis describes work whose ultimate goal is precision mass comparison at the 10^{-11} level. We hope to achieve this precision by measuring the ratio of cyclotron frequencies of individual ions, stored in a Penning trap.

This level of precision should provide interesting results in diverse areas of physics and chemistry:

- The $^3H^+ - {}^3He^+$ mass difference is an important (and currently controversial) parameter in experiments that attempt to measure the electron neutrino rest mass

- The mass difference between $^{14}N^+$ and $^{15}N^+$ ions can be combined with a direct measurement of γ-ray wavelengths to yield a new value of the Avogadro constant.

- Binding energies (and even vibrational excitations) of molecular ions might be measured by weighing the small mass difference $\Delta m = E_{bind} / c^2$ of molecular ions in the trap.

- The trap may be used to study multiply-charged ions as easily as singly-charged ions.

- The traditional applications of mass spectroscopy should benefit from both the accuracy and sensitivity improvement this approach offers over conventional techniques.

Descriptions of several of these experiments were given in a previous dissertation [FLA87] and will not be discussed further here.

Past experiments [VAS81] have shown that ion-ion interactions give shifts in the cyclotron frequency of roughly 10^{-10} per ion, an order of magnitude larger than our desired precision. Commercial ICR spectrometers, limited by both the space charge of larger ion clouds and geometrical effects because of imperfect trapping potentials, have shifts and widths another two orders of magnitude worse. Thus, to reach the 10^{-11} level, we must work with a single ion in a precision trap.

Towards these ends, we have designed a superconducting detector capable of measuring the motion of a single, trapped ion held at 4 K. We have demonstrated that we can, indeed, see a single N_2^+ ion and keep it trapped, isolated, for extended periods of time. Although we have not yet made cyclotron measurements on an isolated ion, we have made preliminary, "crude" cyclotron measurements on small clouds of ions at the 10^{-7} level. These preliminary measurements and, in addition, extensive axial measurements have yielded important calibration data and demonstrate the feasibility of higher precision.

This thesis describes in detail the theory, the apparatus and the techniques we have used to understand and interrogate the behavior of individual ions and small clouds of ions in a Penning trap. The goal of the experiment is to measure cyclotron frequencies; that is, the rate of precession of a charged particle about a magnetic field line. However, to measure this frequency precisely we must keep the particle around long enough to accumulate sufficient phase information. Therefore, we performed all our measurements on the ions in a static-field, quadrupole, crossed E- and B-field trap,

known as a Penning trap. The trapping principle is easy to understand: a quadratic electrostatic potential gives rise to a linear restoring force along the axial direction of the trap; the strong magnetic field provides radial confinement. As we shall discuss at length, this configuration allows us both to trap the ions and detect them at the minor cost of a slightly modified cyclotron frequency.

Chapter II first describes the classical motion of ions held in an ideal Penning trap, then generalizes the discussion to motion in a less than perfect trap. We expect deviations due to the imperfect vacuum, imperfectly shaped trap electrodes, magnetic field drifts and inhomogeneities, and other experimental realities, including special relativity.

To achieve the necessary sensitivity to detect a single ion's motion, we developed a Superconducting Quantum Interference-based detector (SQUID) sensitive enough to measure the tiny current ($< 10^{-14}\ A$ at 160 KHz) a trapped ion produces as it oscillates between the endcaps of the trap. Chapter III describes the theory of RF SQUIDs in sufficient detail to explain the physical principles underlying our detector. This theory helped us combine a commercial RF SQUID with home-made, superconducting resonant circuits and optimize them to make a single-ion detector. The completed detector is described in our Journal of Applied Physics article, included in Chapter III. An appendix to that paper includes our more recent modifications.

The second part of Chapter III provides an introduction to the Fast Fourier

Transform (FFT) and its use in the experiment. Many of our detection techniques rely on analyzing the transient response of the ions to pulsed drives. We analyze these responses in the frequency domain using the FFT. Because the (discrete) FFT can be quite different from the (more familiar) continuous fourier transform, we have included this section to help extend the intuition about "normal" fourier transforms into the discrete-time, sampled-frequency domain where the FFT acts.

Chapter IV describes our apparatus, concentrating on modifications and improvements since the earlier dissertation. The trapping apparatus (as opposed to the magnet and other cryogenic systems) has gone through substantial changes since the earlier work, and we summarize those improvements. We also briefly describe our computer system, and how we used it to mimic several of the more effective bandwidth-narrowing instruments (for example, a lock-in amplifier and a coherent transient averager). We conclude Chapter IV with a description of a typical experiment cycle, going from the disassembled probe to detected ions.

Chapters V and VI interleave theory and experiments, describing predicted and measured results from driving and detecting trapped ions. Chapter V focuses on the axial motion of the ions. We show that the ions can be described by a simple equivalent circuit, and this circuit lets us determine the transient and CW response of the ions and detector when driven by a single, oscillating drive. Predictions include novel effects due to the trap capacitance and strong coupling between the ions and the resonant detector. This one-drive technique has allowed us to determine quickly the

axial resonant frequency of the ions and to make routine measurements of the number of ions in the trap. We then discuss a more sophisticated (and more commonly used) two-drive scheme, in which the ion mixes two independent drives together and oscillates at their sum and difference frequencies. Measurements are more easily interpreted using this two-drive technique because, in almost every case, the coupled ion/detector/trap system behaves in many ways like an isolated ion. Using this technique, we have measured the anharmonicity of the trapping potential, locked the ions to an external frequency source, and performed a Milliken-like counting experiment to demonstrate that we have trapped a single, isolated ion. This final counting experiment, very robust and convincing, was the core experimental goal of this thesis.

Chapter VI continues with the theory and results for the radial motions in the trap. Although these motions are usually orthogonal to the axial motion, we describe how we can couple selectively the radial motions into the axial mode and thus detect them with our SQUID. After introducing the physics of mode-coupling by an inhomogeneous RF electric field, we apply the theory to several experiments: damping the radial motion by coupling it to the (damped) axial detector; measuring the cyclotron frequency by measuring shifts in the *axial* frequency when the inhomogeneous field is resonant between the modes (a "classical" avoided crossing); and measuring the cyclotron resonance directly by simultaneously driving the cyclotron motion and coupling that motion into the axial mode. Because we have just begun

measuring the radial modes, the experimental results in this chapter are more sparse than in Chapter V. However we do present cyclotron measurements using the two techniques as well as strong evidence of radial cooling.

In Chapter VII, we present more results, focusing on measurements of systematic effects which we must understand in order to perform ultra-high precision mass spectroscopy. Thus we include measurements of minimum axial anharmonicities, a characterization of the trap's anti-symmetric potentials (used for shifting the ions within the trap), and techniques that eliminate contaminating ions from the trap. Measured estimates for trap tilt, patches on the trap surfaces, and magnetic field drift are presented, and a theoretical estimate for the magnetic field inhomogeneity is given and a possible measurement suggested. We then outline changes to the existing apparatus that we expect to make in the near-term to improve the signal-to-noise and reproducibility of the experiments. We conclude (the chapter and the thesis) with a somewhat speculative discussion estimating the fundamental (and not-so-fundamental) limits of the experiment for measuring masses to very high precision.

CHAPTER II

ION MOTION IN PERFECT AND IMPERFECT PENNING TRAPS

In this chapter, I will begin to discuss the physics underlying precision mass measurement in Penning traps, focusing in particular on the motion of the trapped ions. After touching briefly on the history of "Penning configuration" ion traps, I will describe the three, orthogonal ion motions in a perfect trap. For massive ions (as opposed to electrons), these motions are described quite well without quantum physics. I then will discuss the implications of the more recent calculations of the electrostatics of *imperfect* Penning traps [GAB83,GAB84,BEA86]. These calculations, driven by the increasingly precise measurements in traps, provide additional information about systematic effects, information that must be used to achieve the increased accuracy that such precision can provide.

Following the discussion of the trap imperfections, I will summarize other perturbations, such as heating by collisions with background gas; the effects of other trapped ions; patch effects; and the effects of magnetic field inhomogeneity, drift, and tilt with the respect to the electrostatic axes. I also discuss relativistic corrections. This chapter concludes with a high-precision prescription to eliminate some of these perturbations that lets us get closer to the final goal: a measurement of the free-space cyclotron frequency, inversely proportional to the ions' mass.

II.A The Ideal Penning Trap

In this section, I will summarize the most basic physics of trapped particles confined in a static-field ("Penning") quadrupole trap. First, I will give a brief history of this type of trap, highlighting its use for high-precision mass measurement. Then, I will describe the three, independent motions of a charged particle in this trap: a harmonic motion in the z-direction, a slightly perturbed cyclotron motion (about the magnetic field lines), and a circular $\vec{E}\times\vec{B}$ (or "magnetron") drift. Finally, I will focus a bit on the z-motion, and demonstrate that detecting this motion will cause it to damp.

II.A.1 History

F. M. Penning is usually credited [DEH67,BRG86] with the idea of trapping charged particles with crossed electric and magnetic fields. In an early paper [PEN36], he observed that at low vapor-pressure, gas in coaxial tubes could sustain glow discharges at much lower voltages when the tube was placed in a strong magnetic field. He attributed this decrease to the greatly increased path-length for the electrons as they followed cycloidal trajectories along the magnetic field lines. (A.W. Hull at G.E. had calculated these trajectories well before Penning's work. [HUL21]) This effect became the basis of the Phillips Vacuum Gauge, later more commonly called "Penning Gauges," which basically work by making a trapped-electron vacuum tube whose voltage-current curve is determined by the background pressure.

After World War II, there were several references to particle traps using this principle. For example, T. R. Pierce devotes an entire chapter in his book on electron

beams [PIE49] to the theory of a crossed-field electron trap. In that same year (1949), the first high-precision measurement of an ion mass in a trap using cyclotron resonance was performed by Hipple, Sommer and Thomas at N. B. S. [HST49,STH51]. They observed protons ejected from their trap when they set an external RF electric field precisely to the protons' cyclotron frequency. Ultimately, they combined this measurement (at the 10^{-5} level) with the proton gyromagnetic ratio to get a precise value for the Faraday, the total charge of a mole of electrons.[1]

In the 1960's, Hans Dehmelt, performing hyperfine spectroscopy experiments on trapped ions in Washington began referring to his DC (or magnetron) traps as "Penning" traps, and the name stuck. Though Penning traps were used extensively for RF spectroscopy in the 60's, they were not resuscitated for fundamental constant work until a single electron was trapped and its g-factor measured precisely, again by the group at Washington [WED73]. Since that time, Penning traps have been used in increasingly precise measurements; for example, the ratio of the electron and positron g-factors [VSD87] and the ratio of the proton and electron masses [VAS81,VMF85].

However, much smaller signals and other technical problems have made it very difficult to reproduce for ions the quantum leap in precision that trapping a single electron made possible. Only very recently [VMF86] were single protons detected;

1. They called their device an "omegatron," a name which, luckily, did not catch on.

this thesis reports the first detection of single, heavy (N_2^+ ions. As yet, no one has reported mass measurements using a single ion.

II.A.2 Summary of Classical Motions

In this section, I will discuss the motions of ions stored in an ideal Penning trap. By "ideal," I mean a trap formed with perfectly machined, infinite electrodes each of which follows surfaces $r^2P_2(\cos\theta) = constant$. (See Fig II.A.2.1) The two electrodes symmetrically above and below the z=0 plane we call "endcaps" and, for trapping, these endcaps are held at equal potentials. The remaining hyperboloid (rotated about the z-axis), we call the "ring." To trap positive ions, the ring must be held at a more negative potential than the endcaps. Since $r^2P_2(\cos\theta)$ satisfies the source-free Poisson equation in radial coordinates [JAC75], we can use the boundary conditions to write down the solution for the electrostatic potential:

$$\Phi(\vec{r}) = V_t \frac{z^2 - \frac{1}{2}\rho^2}{z_0^2 + \frac{1}{2}\rho_0^2} + \Phi_0 \qquad \text{(II.A.2.1)}$$

where V_t is the difference between the endcap and ring potentials, and Φ_0 is an unobservable constant. (The other parameters are defined in Fig II.A.2.1) We see from (II.A.2.1) that the potential is quadratic in the axial direction. That is, the trap provides a harmonic restoring force in the z direction.

To provide radial confinement, we need a constant magnetic field in the z direction:

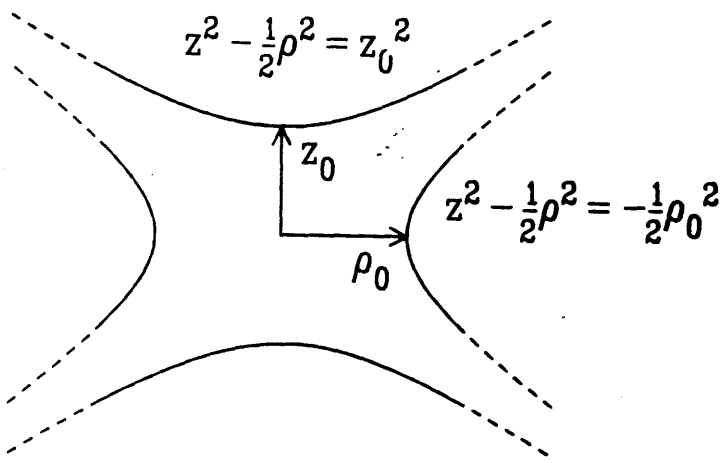

Figure II.A.2.1. The ideal Penning trap, shown in cylindrical coordinates. In our traps, $z_0 = 0.600$ cm and $\rho_0 = 0.696$ cm.

$$\vec{B} = B_0\hat{z} \qquad \text{(II.A.2.2)}$$

For these simple fields, we can write down the equation of motion:

$$\vec{F} = -e\vec{\nabla}\Phi + \frac{e}{c}\vec{v}\times\vec{B} \qquad \text{(II.A.2.3)}$$

and solve it. The solution consists of three, independent harmonic oscillators. Because the solution has been given in several places, ([DEH67, BRG86,FLA87]), I'll only give the results.

In the $\hat{z}$ direction, there are no terms due to the Lorentz force, (since $\vec{B}$ is in the $\hat{z}$ direction) and we get:

$$\ddot{z} + \omega_z^2 z = 0 \qquad \text{(II.A.2.4)}$$

where

$$\omega_z^2 \equiv \frac{eV_t}{md^2}, \qquad \text{and} \qquad \text{(II.A.2.5)}$$

$$d^2 \equiv \tfrac{1}{2} z_0^2 + \tfrac{1}{4}\rho_0^2$$

In our trap, $d^2 = 0.3\text{cm}^2$, and thus to bring ions ions into resonance with our detector at $\omega_z = 10^6\ s^{-1}$ requires a trapping voltage $V_t \approx 0.3$ V per amu. For N_2^+ ions, this voltage is 8 V.

In the radial direction, we get a two-dimensional vector equation:

$$\ddot{\vec{\rho}} - \omega_c\hat{z}\times\dot{\vec{\rho}} - \frac{1}{2}\omega_z^2\vec{\rho} = 0 \qquad \text{(II.A.2.6)}$$

where

$$\omega_c \equiv \frac{eB}{mc} \tag{II.A.2.7}$$

the usual, free-space cyclotron frequency. (In our 8.5×10^4 gauss field, N_2^+ ions precess at 3×10^7 s^{-1}.) The radial motion can be decoupled further into two, independent motions: a slightly perturbed cyclotron motion and an $\vec{E}\times\vec{B}$, or "magnetron" drift. The cyclotron motion arises from the first two terms of (II.A.2.6), modified slightly by the third term. The magnetron motion arises from the second two terms, modified slightly by the first. (The second fact isn't obvious, but can be demonstrated quite easily. Neglecting $\ddot{\vec{\rho}}$, we have:

$$-\omega_c \hat{z}\times\dot{\vec{\rho}} - \frac{1}{2}\omega_z\vec{\rho} = 0 \tag{II.A.2.8}$$

Taking the cross-product of both sides with $\hat{z}$, using $\hat{z}\times(\hat{z}\times\vec{\rho}) = -\vec{\rho}$, yields:

$$\dot{\vec{\rho}} = \frac{1}{2}\frac{\omega_z^2}{\omega_c}\hat{z}\times\vec{\rho} \tag{II.A.2.9}$$

which is the equation for circular motion about the $\hat{z}$-axis with frequency $\omega_m \approx \frac{1}{2}\frac{\omega_z^2}{\omega_c}$. For a given trapping potential, ω_m is roughly independent of the mass. For N_2^+ ions in our magnetic field, resonant with our detector, $\omega_m \approx 1.6\times10^4$ s^{-1}.)

To decouple the two modes in equation II.A.2.6, I follow the discussion of Brown and Gabrielse, who introduce two velocity-like vectors, $\vec{V}^{(\pm)}$:

$$\vec{V}^{(\pm)} = \dot{\vec{\rho}} - \omega_+\hat{z}\times\vec{\rho} \tag{II.A.2.10}$$

where

$$\omega_{\pm} = \frac{1}{2}\left[\omega_c \pm \sqrt{\omega_c^2 - 2\omega_z}\right] \qquad \text{(II.A.2.11)}$$

are the two radial eigenfrequencies.

Using these expressions, Equation II.A.2.6 reduces to

$$\dot{\vec{V}}^{(\pm)} = \omega_{\pm}\hat{z} \times \vec{V}^{(\pm)} \qquad \text{(II.A.2.12)}$$

Thus, the $\vec{V}^{(\pm)}$ are independent, and their motions are quite simple: rotation about $\hat{z}$ with angular frequencies $\omega_{\pm}$. In our experiment, we always operate in the regime $\omega_c \gg \omega_z$. In that case, we can simplify (II.A.2.11) by expanding the square root to get:

$$\omega_- = \frac{1}{2}\frac{\omega_z^2}{\omega_c}\left(1 + \frac{1}{2}\frac{\omega_z^2}{\omega_c^2} + \cdots\right) \qquad \text{(II.A.2.13)}$$
$$\omega_+ = \omega_c - \omega_-$$

Thus ω_- corresponds to magnetron motion and ω_+ corresponds to cyclotron motion, each perturbed slightly from the naive values given above in Equations II.A.2.7 and II.A.2.9.

We see from the definition of $\vec{V}^{(\pm)}$ that the radial position and velocity correspond almost independently to magnetron and cyclotron motion, respectively. That is, since $\omega_- \ll \omega_+$, we can use (II.A.2.10) to show that $\dot{\vec{\rho}} \approx \vec{V}^{(+)}$ and $\vec{\rho} \approx \frac{1}{\omega_+}\hat{z} \times \vec{V}^{(-)}$. To understand further the physical meaning of the $\vec{V}^{(\pm)}$ vectors, we can solve for the radial velocity:

$$\dot{\vec{\rho}} = \frac{\omega_+\vec{V}^{(+)} - \omega_-\vec{V}^{(-)}}{\omega_+ - \omega_-} \qquad \text{(II.A.2.14)}$$

and consider what happens to $\dot{\vec{\rho}}$ when either $\vec{V}^{(+)}$ or $\vec{V}^{(-)}$ vanishes. Using $\omega_+ \gg \omega_-$, for

$\vec{V}^{(-)}=0$, we see that $\vec{V}^{(+)} \approx \dot{\vec{\rho}}$ and thus $\vec{V}^{(+)}$ is simply the velocity of the cyclotron motion. We therefore can define $\rho_c \equiv \frac{V^{(+)}}{\omega_+}$ to be the cyclotron radius. On the other hand, for $\vec{V}^{(+)}=0$, the magnetron vector, $\vec{V}^{(-)} \approx -\frac{\omega_+}{\omega_-}\dot{\vec{\rho}}$, and thus $\vec{V}^{(-)}$ is much larger than magnetron orbit velocity; that is, $\rho_m \equiv \frac{V^{(-)}}{\omega_+}$ and *not* $\frac{V^{(-)}}{\omega_-}$.

Brown and Gabrielse [BRG86] give a nice expression for the energy in the two radial modes:

$$H_\rho = \frac{1}{2}m\left[\frac{\omega_+\vec{V}^{(+)^2} - \omega_-\vec{V}^{(-)^2}}{\omega_+ - \omega_-}\right] \qquad \text{(II.A.2.15)}$$

From the previous discussion, and with the same assumptions, we see that the first term is $H_{cyc} \approx \frac{1}{2}m\vec{V}^{(+)^2}$. Since $\vec{V}^{(+)}$ is the cyclotron velocity, we see that the cyclotron contribution to the energy is the kinetic energy of the orbiting particle. The second term, on the other hand, is $H_{mag} \approx -\frac{1}{2}m\frac{\omega_-}{\omega_+}\vec{V}^{(-)^2}$ Using $\rho_m \approx \frac{V^{(-)}}{\omega_+}$ and $\omega_-\omega_+ = \frac{1}{2}\omega_z^2$, we find that $H_{mag} = -\frac{1}{4}m\omega_z^2\rho_m^2$. Thus the magnetron orbit contributes only potential energy: its kinetic energy is very much smaller.

The sign of the radial potential energy causes some authors great concern, and invites discussion about whether the Penning trap is a true "trap." While it is true that the radial motion is unstable, and that collisions tend to cause ions to diffuse slowly

out, under typical experimental vacuum, we expect ions to stay "trapped" for periods of roughly one month. (See the discussion of collisions later in this chapter)

Finally, we might ask whether the classical description is sufficient to describe these harmonic oscillators at liquid Helium temperatures. We can answer that question by calculating the average number of quanta associated with a harmonic oscillator coupled to a thermal bath at temperature T:

$$N \approx \frac{kT}{h\nu} \qquad \text{(II.A.2.16)}$$

A single, trapped N_2^+ ion, for example, has a cyclotron frequency, in our magnetic field of 85000 gauss, of 4.6 MHz. At this frequency, we find, at 4.2 K, $N \approx 2\times10^4$. Thus we can conclude, quite safely, that the classical description will suffice. Since the cyclotron motion has the highest frequency of the normal modes, it must have the lowest N. Therefore we can conclude that the classical equations of motion suffice to describe *all* the motions in the trap. (Even if they did not, the quantum mechanical description does not hold any real surprises, anyway. [BRG86])

II.A.3 Damping Because of Detection

There is no inherent damping in the equations just discussed, (II.A.2.4) and (II.A.2.6). We can safely neglect the radiation damping of the cyclotron motion: for example, at 4.6 MHz, N_2^+ ions have a free-space damping time of 160 000 years! (In fact, even *this* time would substantially underestimate the damping because the trap essentially forms a cavity with dimensions far smaller than the wavelength. Thus there

are no modes into which the cyclotron can decay. [GAD85]) However, when we hook the ions to a detector, we immediately introduce damping. Wineland and Dehmelt [WID75] used the work-energy theorem to calculate the axial damping when a dissipative load is connected across the trap. Rather than repeat their argument, for variety let me discuss the same relationship using Green's reciprocity theorem. [JAC75]

The reciprocity theorem states that, for any given geometry, two solutions to Maxwell's equations are related by:

$$\int dV\, \rho\, \phi' + \int dA\, \sigma\, \phi' = \int dV\, \rho'\, \phi + \int dA\, \sigma'\, \phi \qquad \text{(II.A.3.1)}$$

where ϕ, ρ and σ are, respectively, the potential, charge density and surface charge density for one solution and ϕ', ρ' and σ' for the second.

We can apply this theorem to a particle within a parallel plate capacitor of spacing $2z_0$. (See fig II.A.3.1) In this case, the un-primed system corresponds to the particle within the capacitor, and the primed systems corresponds to the potentials on the surfaces. Since the charge density, $\rho(\vec{z})$ is just a delta function of the particle's position, $\hat{z}$, the first term of (II.A.3.1) yields $e\phi'(\hat{z})$. Because the upper plate of the capacitor is an equipotential, the surface integral yields $V \int dA\, \sigma \equiv V \cdot Q$. The right side is zero: there are no charges in the primed system, and ϕ vanishes on the surfaces. Thus, quite generally (since we haven't yet used the specific geometry of the parallel plate capacitor system):

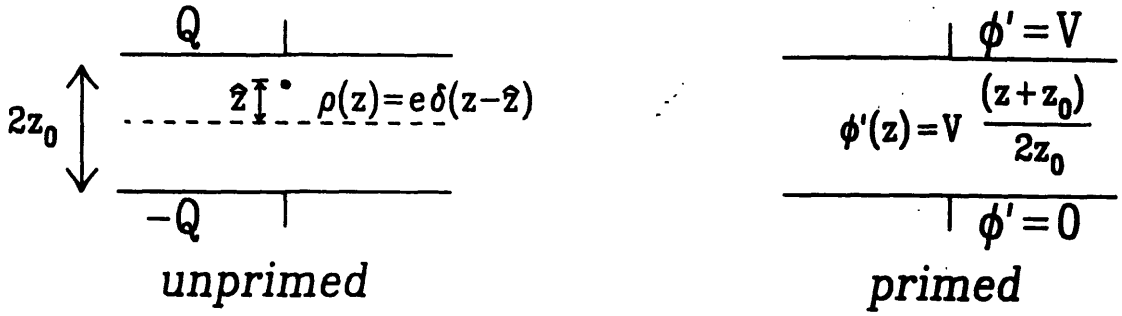

Figure II.A.3.1. Charge induced by an ion moving in a parallel plate capacitor. (This geometry is for the reciprocity theorem in the text.)

$$Q = -\frac{e\phi'(\hat{z})}{V} \qquad \text{(II.A.3.2)}$$

and, in the case of a parallel-plate capacitor:

$$Q = -\frac{e}{2}\left[\frac{\hat{z}}{z_0} + 1\right] \qquad \text{(II.A.3.3)}$$

At this point, I drop the $\hat{z}$ and just write z, since there can be no confusion. Since $i = \dot{Q}$, we have $i = -\frac{e\dot{z}}{2z_0}$.

In fact, the Penning trap is not a parallel plate capacitor, but to first order we can follow Gabrielse and correct this expression:

$$i = -\frac{eB_1}{2z_0}\dot{z} \qquad \text{(II.A.3.4)}$$

where B_1 is constant of order unity computed for various traps in [GAB84], predicted (and approximately measured) in our trap to be 0.8.

Now let us add a detector to measure the axial motion. Assume, for the moment, that we use some circuit elements (coils, capacitors, resistors) to detect the motion. We will show that the real part of the impedance of this detector will cause the ions to lose energy. (The imaginary part causes frequency shifts, as I will discuss in Chapter V.)

As the ion oscillates in the trap, the current (II.A.3.4) it induces, in turn, causes a voltage drop across the detector. The part of the voltage associated with the real part of the detector's impedance is in phase with the current: $V = \mathrm{Re}\,Z\,\frac{eB_1}{2z_0}\dot{z}$. (see

Fig II.A.3.2) This voltage creates an electric field in the trap which acts back on the ions: $F = -eB_1 \frac{\dot{V}}{2z_0} \dot{z}$. (This B_1 is the same constant as in Equation II.A.3.4.) So now the ion's equation of motion, (II.A.2.4) reads:

$$\ddot{z} + \omega_z^2 z = -\left\{ \frac{1}{m} \left[\frac{eB_1}{2z_0} \right]^2 \mathrm{Re}Z \right\} \dot{z} \qquad \text{(II.A.3.5)}$$

We can write this equation simply:

$$\ddot{z} + \gamma_z \dot{z} + \omega_z^2 z = 0 \qquad \text{(II.A.3.6)}$$

where

$$\gamma_z \equiv \frac{1}{m} \left[\frac{eB_1}{2z_0} \right]^2 \mathrm{Re}Z \qquad \text{(II.A.3.7)}$$

For our detector, which is a resonant tuned circuit with an inductance L and quality factor Q, on resonance has a maximum $\mathrm{Re}Z = \omega_z L Q \approx 1.5 \times 10^8\ \Omega$. For one N_2^+ ion, then, $\gamma_z \approx 0.3\ s^{-1}$. The imaginary part of the detector's impedance (that is, terms like $i\gamma_z \dot{z}$) do not do not cause damping. Instead, it shifts the phase of the induced voltage relative to the current, and hence relative to the ions, and thus shifts the effective resonance frequency.

For more than one ion in a perfect trap, we shall see that the center-of-mass motion is identical to a single ion's motion. However the damping increases linearly with the number of ions. We can explain this increased damping quite simply: N ions will induce N times the charge in the upper endcap and thus N times larger current in the detector. (Since these ions are moving as one, the

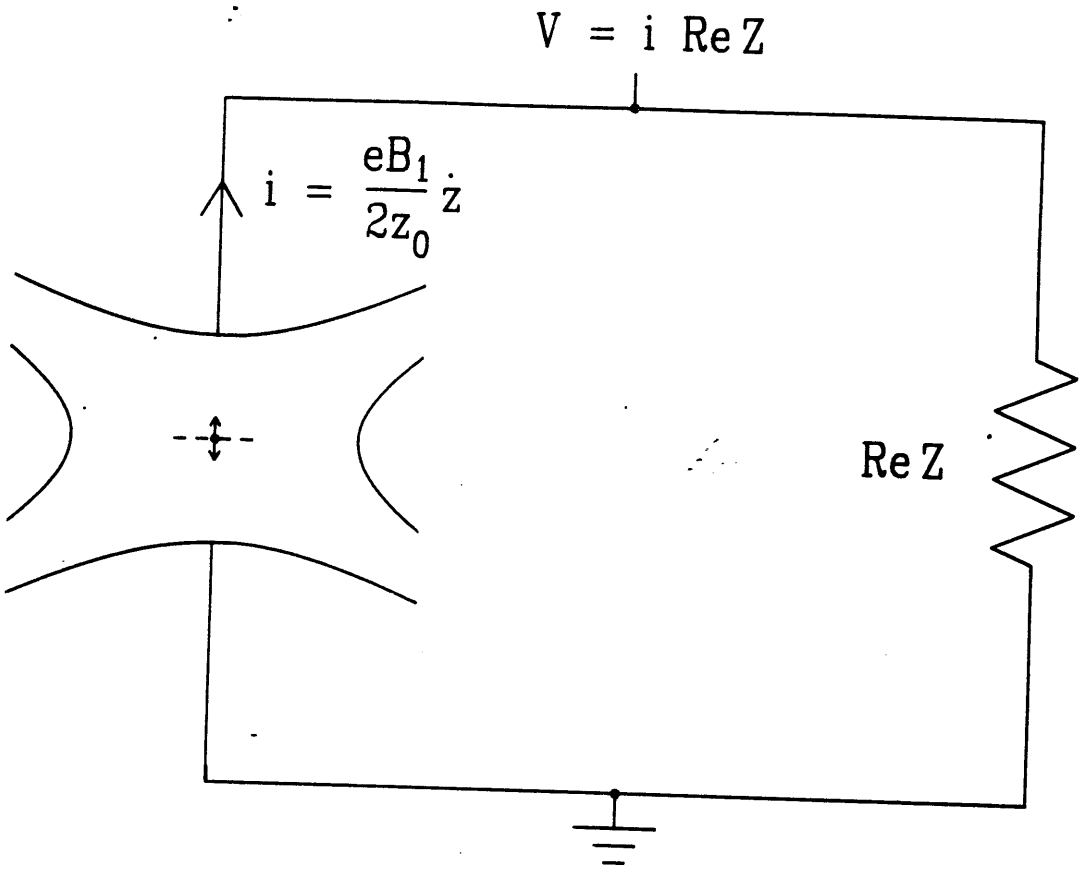

Figure II.A.3.2. Detector causes damping. The ion induces a current which causes a voltage drop across the detector. This voltage acts back on the ion, causing the ion to damp.

individual induced currents are all in phase and simply add.) The voltage drop induced at the detector thus will be N times the voltage induced by a single ion. This larger voltage then creates a larger electric field which affects each ion to an N times greater extent. Therefore, the damping term (the right hand side of II.A.3.5) will by N times larger per ion. This dependence of γ_z on ion number will be crucial for several ion counting schemes.

The proportionality of γ_z and N also has a somewhat counter-intuitive side effect: the total induced current due to an exactly resonant, harmonic drive will be independent of the number of ions in the trap. Since the response of a harmonic oscillator driven on its resonance depends on γ_z^{-1}, the response of *each* ion will decrease with the number of ions. However, since the total current induced is the sum of these individual contributions, the total signal, on resonance, will be the same, regardless of the number of ions in the trap.

II.B Real Axial Motion

In the last section, I described the motion of ions in an ideal Penning trap. In this section, however, I will begin to detail the motion of ions in a *real* Penning trap, in particular, discussing the axial motion of ions in an imperfect trap. I will first describe driving and shifting the ions within the trap using axially anti-symmetric potentials. Then, I will discuss the symmetric trap imperfections, focusing on effects of a term z^4 in the trap potential. Finally, the section ends with a discussion about

compensating the trap to minimize the effects of these anharmonicities.

First, let us try and write down a more complete potential than (II.A.2.1). At this stage, though, we will continue to assume that the trap is axially symmetric. (deviations from this assumption are discussed in section II.C.7.) We can use this symmetry to write down the most general potential that satisfies Maxwell's equations in spherical coordinates [JAC75]:

$$\Phi(\vec{r}) = \sum_{l=0}^{\infty} A_l \, r^l \, P_l(\cos\theta) \qquad \text{(II.B.1)}$$

where the P_l are Legendre polynomials. The terms with l even are even under $z \rightarrow -z$, while those terms with l odd are odd.

Three potentials are under the control of the experimenters: V_r, the potential on the ring, and $V_{\pm}$, the potentials on the upper and lower endcaps. A useful way to combine these potentials with (II.B.1) is to separate the sum into even and odd terms. These separate potentials must therefore arise independently from the even and odd boundary conditions established by V_r and $V_{\pm}$. (See Figure II.B.1) The even boundary condition we will associate with trapping the ions. That boundary condition gives rise to a potential:

$$\Phi_s = \frac{1}{2} \sum_{k \text{ even}} C_k \left[\frac{r}{d}\right]^k P_k(\cos\theta) \qquad \text{(II.B.2)}$$

(where d is the trap size, given by Equation II.A.2.5). In this expression, the leading term is identical to Equation II.A.2.1, and thus provides the harmonic trapping in the

z-direction.

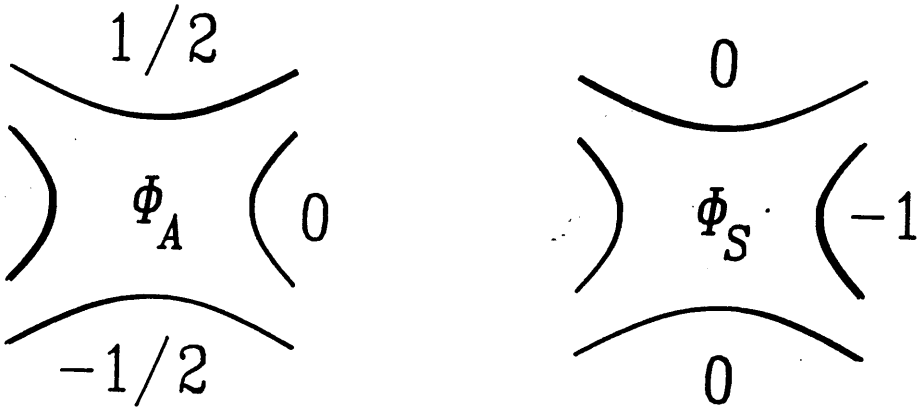

Figure II.B.1. Boundary condition for the symmetric potential, Φ_S, and the anti-symmetric potential, Φ_A.

The odd boundary condition gives rise to the anti-symmetric potential:

$$\Phi_A = \frac{1}{2} \sum_{k\,\text{odd}} B_k \left[\frac{r}{z_0}\right]^k P_k(\cos\theta) \qquad \text{(II.B.3)}$$

The first term, $\Phi_A^{(1)} = \frac{B_1 z}{2z_0}$ yields a constant electric field, like a parallel-plate capacitor, and thus Φ_A can be used to shift, and therefore drive, the ions trapped by Φ_s.

We can express the amount of trapping and shifting due to arbitrary endcap potentials by dividing those potentials into symmetric and anti-symmetric parts:

$$\Phi = (V_+ - V_-)\Phi_A + \left[\frac{V_+ + V_-}{2} - V_r\right]\Phi_s + C \qquad \text{(II.B.4)}$$

where C is an unobservable constant. We will now examine more carefully the effects of these higher order terms in (II.B.3) and (II.B.2).

Anti-Symmetric Part

Anti-symmetric potentials can drive the ions. We used this fact implicitly in the last section to determine the damping: the ions are, in effect, driven by the voltage they induce in the detector, and that drive damps the ions. Using (II.B.3) above, we can determine, to the next highest order, the effects of an anti-symmetric potential. Writing out the first two terms of (II.B.3), we get:

$$\Phi_A = B_1 \frac{z}{2z_0} + B_3 \frac{z^3 - \frac{3}{2} z\rho^2}{2z_0^3} + \cdots \qquad \text{(II.B.5)}$$

Therefore a voltage V_A applied across the endcaps of the trap leads to an axial force on the ions:

$$m\ddot{z} = -\frac{eV_A}{2z_0}\left[B_1 - \frac{3}{2}B_3\frac{\rho^2}{z_0^2} + 3B_3\frac{z^2}{z_0^2}\right] \qquad \text{(II.B.6)}$$

Thus we see two effects due to B_3. First, the applied force tends to decrease as the trapping radius increases. In addition, then, using the reciprocity theorem from the last section, ions with a large magnetron orbit (thus a large, constant ρ) will induce less current in the endcaps. This effect makes some sense in light of the physical deviations from the parallel capacitor model of the last section. At large radii, not only does the plate separation appear larger, but field lines also begin to head off to the ring.

The second effect due to B_3 is that the applied force is non-linear in z, and, in fact, increases at larger orbits. As I will now discuss, a constant voltage applied across the trap therefore can change the axial resonance frequency when the equilibrium position changes. In addition, we can throw ions out of the trap by displacing the center of the ion cloud into the lower endcaps. As we shall see, the voltage required to remove the ions depends on B_3.

To analyze the effects of the B_3 term, let us consider the motion, in general, of a simple harmonic oscillator with an additional constant and quadratic force term:

$$\ddot{z} + \omega_z^2 z + a + b z^2 = 0 \qquad \text{(II.B.7)}$$

The harmonic motion, for small amplitude oscillations, will be around some equilibrium point, $\hat{z}$, at which the constant term effectively vanishes. (This shift of the equilibrium position is like the shift for a mass on a spring in a gravitational field.) As we shall see, one effect of the quadratic term is to shift the effective spring constant with $\hat{z}$.

To keep contact with the ion trap problem, the constants a and b in Equation II.B.7 are:

$$a = \frac{eB_1}{2mz_0} V_A \qquad \text{(II.B.8)}$$

$$b = \frac{3eB_3}{2mz_0^3} V_A$$

Recall, too, that ω_z^2 comes from the symmetric part of the potential. At this stage, we only consider the dominant term of that potential.

We can find both the equilibrium position and the oscillation frequency, ω_{eff}, around that equilibrium from (II.B.7). We make the transformation $z \rightarrow (z - \hat{z})$, and seek the $\hat{z}$ that eliminates the constant term a. The algebra is simple:

$$b \hat{z}^2 + \omega_z^2 \hat{z} + a = 0 \qquad \text{(II.B.9)}$$

At that value of $\hat{z}$, the term in the transformed equation of motion that is linear in z will be:

$$\omega_{eff}^2 = \omega_z^2 + 2b\hat{z} \qquad \text{(II.B.10)}$$

When $b = 0$ (a parallel plate capacitor, for example), the resonant frequency does not

shift. However, when $b \neq 0$, we must solve the quadratic equation, (II.B.9). Discarding the unphysical solution, we get:

$$\hat{z} = \frac{\sqrt{\omega_z^4 - 4ab} - \omega_z^2}{2b} \tag{II.B.11}$$

and

$$\omega_{eff}^2 = \sqrt{\omega_z^4 - 4ab}$$

There are two limits which interest us. First, when a and b are both small (corresponding to a small anti-symmetric potential), we find:

$$\hat{z} \approx -\frac{a}{\omega_z^2} = -\frac{B_1 d^2}{2z_0^2}\frac{V_A}{V_{trap}} \tag{II.B.12}$$

and

$$\omega_{eff} \approx \omega_z - \frac{ab}{\omega_z^3} = \omega_z \left[1 - \frac{3d^4}{4z_0^4} B_1 B_3 \frac{V_A^2}{V_{trap}^2}\right]$$

Thus, the shifted position remains linear with applied voltage (as if $B_3 = 0$), and we get a small quadratic shift in the resonant frequency proportional to $B_1 B_3$. Using this shift, we can measure $B_1 B_3$ for a real trap. (See Chapter VII)

The other interesting limit of Equation II.B.11 is the value for which $\hat{z} = \pm z_0$: the value at which we force the ion cloud into the upper or lower endcap. Since both a and b depend on V_A, I return at this point to the trap variables, and restate the question: By changing only the lower endcap voltage, V_-, can we drive the ions out the trap? For simplicity, let $V_+ = 0$. Then, in Equation II.B.7, ω_z^2 comes from the symmetric part of the potential, $(V_-/2 - V_r)$, and a and b come from the anti-symmetric part, $-V_-$. Thus, using (II.B.4) and (II.A.2.5), we can solve (II.B.9) for

$\hat{z} = \pm z_0$:

$$V_{\mp}^{kill} = \frac{2}{1 \mp \frac{d^2}{z_0^2}(B_1 + 3B_3)} V_r \qquad \text{(II.B.13)}$$

where $V_{\mp}^{kill}$ is the voltage required on the lower cap to "kill" the ions by bashing them into the endcaps.

Besides being useful as a method to expell ions from the trap, $V_{\mp}^{kill}$ provides a method for measuring B_1 and B_3. First, we can measure the product, $B_1 B_3$ using (II.B.12), by measuring the quadratic shift in the resonance frequency, written in terms of trap parameters:

$$\frac{\Delta\omega_z}{\omega_z} = \frac{3}{4}\frac{d^4}{z_0^4} B_1 B_3 \frac{V_{\mp}^2}{V_r^2} \qquad \text{(II.B.14)}$$

Then, by measuring the voltage $V_{\mp}^{kill}$ at which the ions hit the lower endcap, we get an estimate of the sum $B_1 + 3B_3$. In practice, however, we've found it difficult to make a precise, consistent measurement of $V_{\mp}^{kill}$, and thus must be satisfied with limits on B_1 and B_3. (See Chapter VII)

Anharmonicity

Let us return now to the effects of higher order terms in the trapping potential, Φ_s. In particular, I will discuss the non-linear response of strongly-driven ions; why we'd rather keep their response linear; and, therefore, what we've done to help

eliminate the non-linear effects.

The first few terms of Equation II.B.2 are:

$$\Phi_s = \frac{1}{2}C_2\frac{z^2 - 1/2\rho^2}{d^2} + \frac{1}{2}C_4\frac{z^4 - 3z^2\rho^2 + 3/8\rho^4}{d^4} + \cdots \qquad \text{(II.B.15)}$$

The first term represents the "ideal" penning trap. In what follows, I will assume $C_2 = 1$, which is almost true, and can be made exactly true by a minor renormalization of d. Using (II.B.15), we can write down the equation of motion in the z direction for a trapping potential $V_T\Phi_s$:

$$\ddot{z} + \omega_z^2(1 - 3C_4\frac{\rho^2}{d^2} + 2C_4\frac{z^2}{d^2})z = 0 \qquad \text{(II.B.16)}$$

where

$$\omega_z^2 = \frac{eV_T}{md^2} \qquad \text{(II.B.17)}$$

as before. The two additional terms in (II.B.16), like the additional terms in (II.B.6) express, first, a radius-dependent shift in ω_z, and second, a non-linear term in the potential.

For our purposes, the dominant effect of the the non-linear term is that it makes the ions' resonant frequency amplitude dependent. Landau and Lifshitz [LAL76] give a particularly elegant and compact treatment of the non-linear oscillator, and I'll sketch their results here.

They expand the solution $z(t)$ in a successive approximation series, starting from $z = a\cos\omega t$, treating the z^3 term as a perturbation. In the resulting perturbations

series, terms that look like resonance driving terms arise. For example, since $\cos^3 \omega t = (\cos 3\omega t + 3\cos\omega t)/4$, the $\cos\omega t$ part can look, to the ions, like a resonant driving term. Such terms are clearly unphysical, and ω must be shifted slightly to make these terms vanish. Using this technique, Landau and Lifshitz show that the frequency shift, to second order in a (the ion's amplitude at ω), will be given by:

$$\frac{\Delta\omega}{\omega} = \frac{3}{4} C_4 \frac{a^2}{d^2} \qquad \text{(II.B.18)}$$

Marion [MAR70] uses the same technique in a somewhat expanded format and obtains the identical result.

Small though this correction may seem, it can never-the-less have dramatic effects when this frequency shift is comparable to the width of the resonance. Most significantly, the resonance will become "hysteretic;" that is, the resulting amplitude from an external drive will depend on the recent history of the ion.

Landau and Lifshitz's explanation of hysteresis uses an ingeniously simple argument. If the dominant effect of the anharmonicity is to shift the resonance, then, as an ion responds to an external drive, V_d, this frequency shift will begin to push the ions nearer or further from resonance, and hence, change the response. Therefore, we must solve self-consistently for the driven, steady-state response. In the narrow resonance approximation, we get:

$$z = -\frac{eB_1}{4mz_0\omega_z} \frac{V_d}{(\omega_d - \omega(z)) + i\gamma_z/2} \qquad \text{(II.B.19)}$$

where, from (II.B.18):

$$\omega(z) = \omega_z \left(1 - \frac{3}{4} C_4 \frac{z^2}{d^2}\right) \qquad \text{(II.B.20)}$$

The resulting cubic equation (in z) can have either one or three real roots, depending on the relative sizes of V_d and γ_z. Following usual analytic techniques [e.g., ABS70], we find that, when the drive exceeds a critical value, V_d^{crit}:

$$V_d^{crit} = \left[\frac{16}{3} \frac{1}{C_4} \frac{\gamma_z^3}{\omega_z^3} \frac{z_0^2}{d^2} \frac{1}{B_1^2}\right]^{\frac{1}{2}} V_t \qquad \text{(II.B.21)}$$

then Equation II.B.19 above will have, at some detuning, three real roots. Note also that V_d^{crit} is the drive that shifts, at maximum response, the resonant frequency, (II.B.18), by γ_z. (See Figure II.B.2a) Careful analysis [LAL76] shows that the intermediate root always corresponds to an unstable response. Therefore, were we to sweep the drive frequency from left to right (in Figure II.B.2b), the response would build up until it reached the critical point A. At that point, the driven response catastrophically drops to point B and thereafter follows the lorentzian tail. Sweeping back the other way, the response rises until it reaches point C, where, to avoid the unstable branch, jumps up to point D.

There are several interesting features of these non-linear, hysteretic resonances. First, anharmonicity does not limit the absolute, attainable peak response; it merely changes the required drive frequency and makes this peak accessible only when swept in the proper direction. That is, at least for one direction of sweep (the one that "pushes away" the ions from the drive), the absolute maximum response remains:

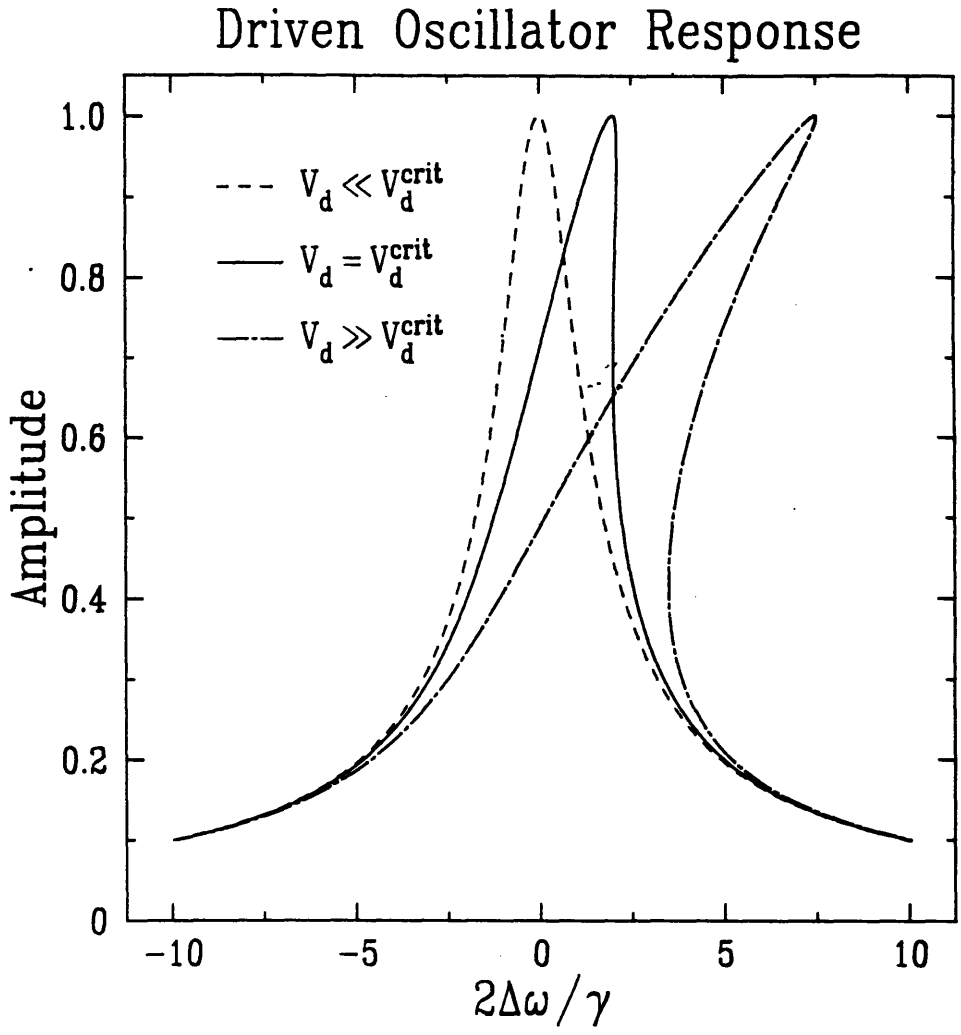

Figure II.B.2a. Anharmonic oscillator response for three different levels of anharmonicity. Between each graph, C_4 was changed by a factor of four.

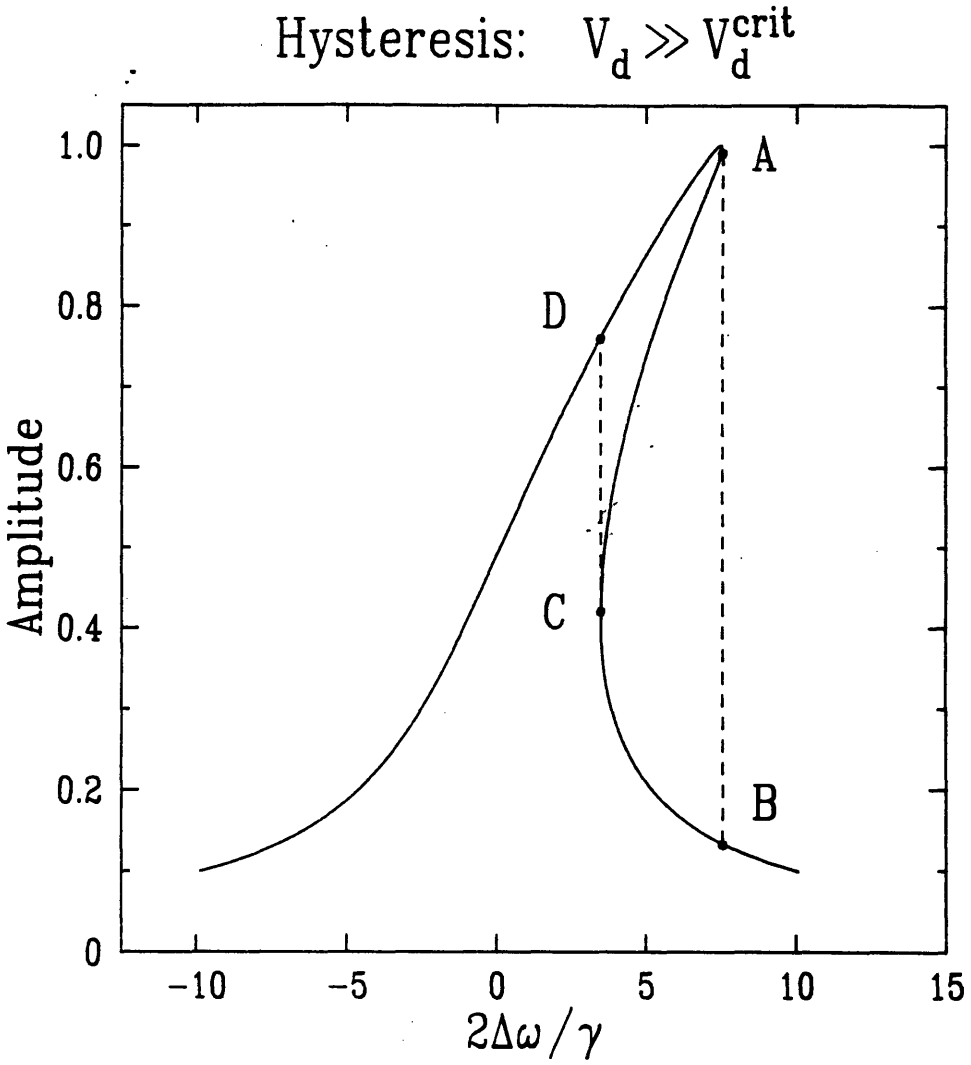

Figure II.B.2b. Hysteretic response. When swept left-to-right, the ion follows the upper curve to point A, then drops precipitously to point B. Sweeping back right-to-left, the ion follows the lower curve past point B all the way to point C, then leaps up to point D. (The underlying curve is the same as the most anharmonic curve of Figure II.B.2a.)

$$z_{peak} = \frac{eB_1}{2mz_0\omega_z\gamma_z} V_{drive} \qquad \text{(II.B.22)}$$

However, even though the *peak* response may not decline, there are still good reasons to try to remain in the linear regime. For example, after sweeping across the catastrophic decrease in signal, experimentally we have found that a great deal of energy is left in the ions, though no longer in phase with the drive. Also, the phase of the response at the peak makes it impossible to lock the ions to an external frequency source. In addition, when the ions are excited by a sharp pulse, their non-linear response becomes strongly dependent on the initial conditions, and thus high precision schemes (like separated oscillatory fields) become difficult, if not impossible.

Therefore, even though the maximum possible ion signal does not necessarily decrease because of non-linearity, we still would prefer to drive at an amplitude less that V_d^{crit}. However, this requirement becomes harder and harder to satisfy as we decrease the number of ions. We can see that this effect is contained in Equation II.B.21. Since the critical drive depends on $\gamma_z^{3/2}$, and γ_z depends linearly on the number of ions, the critical drive must decrease dramatically with the number of trapped ions. Assuming that we want to maximize their response, and yet retain linearity, we must therefore drive the ions with a drive a bit smaller than V_d^{crit}. Since the peak oscillator response varies like γ_z^{-1}, each ion, overall, will thus have a peak response that goes like $\gamma_z^{1/2}$. However, when fewer ions are in the trap, the total induced current decreases, too. Therefore, in summary, the total detector response

possible for a non-hysteretic resonance will increase like $\gamma_z^{3/2}$, and thus decrease strongly when there are fewer trapped ions. Currently, to detect "reasonably" by linear resonance a single N_2^+ ion requires $C_4 < 2\times10^{-5}$. Thus a major hurdle in detecting small numbers of ions in the trap is to overcome this anharmonicity problem.

Compensation

In order to improve the harmonicity of the trap, then, a set of compensations electrodes, placed between the ring and endcap can be added [VWE76]. Although they had been used for several years prior, the first numerical analysis of the fields produced by these extra "guard" rings (a relaxation calculation by Gabrielse [GAB83]) produced reasonable estimates for the effectiveness of such guard rings. In particular, Gabrielse calculated the change in C_4 produced by a given change in the potential on the guard ring. He also showed that, for traps constructed at that time, these changes in C_4 always were accompanied by large shifts in C_2; that is, minor guard ring adjustments shifted the resonant frequency, too. Thus, tuning the trap could be quite difficult [BRG86].

However, Gabrielse's work (and confirming work by Beaty [BEA86]) pointed to a trap construction that minimized these concurrent shifts of the resonant frequency. In particular, they suggested that for $\rho_0 = 1.16 z_0$, no shift in the resonant frequency would take place. (Although the surfaces of a Penning trap must follow hyperboloids with fixed values of $z^2 - \rho^2/2$, the asymptotes, and thus the ratio of z_0 to ρ_0 remains

arbitrary.) We followed that "optimal" prescription in our traps. The degree to which the guard rings actually were decoupled from the resonant frequency, as well as their effectiveness in canceling out anharmonicities will be discussed further in Chapter VII.

II.C Other Perturbations

In addition to the trap electrostatic perturbations just discussed, there are two other classes of deviation from ideality. The first kind, the most troublesome for us (and more generally, it appears [WIN87,GAB87,MOO87]) come from contamination both by background gas and, especially other species of trapped ions. The second class contains a whole set of field perturbations: magnetic bottle shifts, magnetic field drift, misalignment between the magnetic and electrostatic axes, stray electric fields because of patch effects, and azimuthal asymmetry in the endcaps. The effects of these field perturbations must be anticipated for high precision measurements.

In this section, then, I will address both classes of perturbation. Contamination results in inescapable instability, broadening or shifts, and the only solution appears to be to avoid it altogether. For most of the field perturbations, the effects can be measured and several measurements can be used to extrapolate to zero perturbation. For the final few, clever prescriptions have been developed [BRG82] that eliminate these effects entirely from a final, computed value of the free-space cyclotron frequency.

II.C.1 Collisions with Neutrals

Collisions between trapped ions and background neutral atoms have two detrimental features, one annoying, one quite troublesome. First, random, disorienting collisions with thermal background atoms tends to cause the magnetron orbit of the trapped ions to increase, and, eventually, will force the ions out of the trap. Second, since these collisions can interrupt the phase of driven motion, background gas can broaden the resonance in precision cyclotron measurements, and, in fact, can be the ultimate limit to the measurement. In this section, I will discuss radial diffusion first, and then estimate the limit on cyclotron precision due to collisional broadening.

We can use a very naive model to illustrate radial diffusion due to collisions. The purpose of the calculation is simply to give an order-of-magnitude estimate for the time scale on which such diffusion might occur. As we shall see, the final result is that the ions perform a radial random walk, with steps $\Delta\rho \approx \frac{\omega_m}{\omega_c}\rho$ ever increasing in size.

Let us assume the only effect of a collision is to randomly re-orient the ion's velocity. In addition, we assume that the z-motion is either damped (because is is connected to a detector) or driven. In this case, we can ignore that motion entirely. While this is not a particularly realistic model, we certainly can use it to obtain the kind of estimate we want.

As discussed in Section II.A.2, the radial velocity of the ion is contained almost entirely in the cyclotron velocity, $\omega_c \rho_c$, while the radial position is given, to the same

accuracy, by the magnetron orbit, ρ_m. Using complex notation to write $\vec{p} = x + iy$, we can express the radial position before a collision as:

$$\rho_{before} = \rho_m e^{i\theta_m} + \rho_c e^{i\theta_c} \qquad \text{(II.C.1.1)}$$

where θ_m and θ_c are the phases of the motion. Since we have assumed that the collision only reorients the velocity, not its magnitude, ρ_c will be unchanged by the collision:

$$\rho_{after} = \rho_m' e^{i\theta_m'} + \rho_c e^{i\theta_c'} \qquad \text{(II.C.1.2)}$$

Because the position itself remains unchanged, we must have $\rho_{before} = \rho_{after}$, and thus:

$$\rho_m' e^{i\theta_m'} = \rho_m e^{i\theta_m} + \rho_c e^{i\theta_c}(1 - e^{i\Delta\theta}) \qquad \text{(II.C.1.3)}$$

where $\Delta\theta$ is the phase change because of the collision. Taking the magnitude of both side of (II.C.1.3) and averaging over the *random* $\Delta\theta$ yields:

$$\langle \rho_m'^2 \rangle_{\Delta\theta} = \rho_m^2 + 2\rho_c^2 - 2\rho_m \rho_c \cos(\theta_m - \theta_c) \qquad \text{(II.C.1.4)}$$

(For more realistic collisions, we would still perform this average. Though some numerical factors might change, the structure of the result will remain the same.) Assuming that the relative phases of θ_m and θ_c are unimportant—for example, by asserting that the background gas is uniformly distributed throughout the trap—we can neglect the third term in (II.C.1.4) and we obtain the simple estimate:

$$\langle \rho_m'^2 \rangle = \rho_m^2 + 2\rho_c^2 \qquad \text{(II.C.1.5)}$$

Thus, on average, each collision tends to *increase* the magnetron orbit size.

If we assume the background gas remains in thermal equilibrium (that is, the trapped ion is a minor perturbation that does not significantly heat up the background) then it is reasonable to assume that, in the long run, the harmonic oscillators (the three trap modes) will have their energies equally partitioned, and thus $E_{mag} \approx E_{cyc}$. In this case, we can relate the cyclotron and magnetron orbit sizes (using Equation II.A.2.15): $\rho_c{}^2 = \frac{\omega_m}{\omega_c}\rho_m{}^2$, and, per collision, we have:

$$\Delta\rho_m \approx \frac{\omega_m}{\omega_c}\rho_m \qquad \text{(II.C.1.6)}$$

Thus the ions diffuse away with geometrically increasing step size.

To estimate the time scale for the diffusion, we must incorporate the mean time between scattering events, $\tau = (n\sigma v)^{-1}$. Note that v (as usual, dominated by the cyclotron motion) increases with increasing orbit size:

$$v = \omega_c \rho_c \approx \omega_z \rho_m \qquad \text{(II.C.1.7)}$$

If we assume, on average, an increase of $\Delta\rho_m$ every $\frac{1}{n\sigma\omega_z\rho_m}$ seconds, we get:

$$\dot{\rho}_m = n\sigma\frac{\omega_z\omega_m}{\omega_c}\rho^2 \qquad \text{(II.C.1.8)}$$

and thus we see that the diffusion grows quite quickly with increasing orbit size. Therefore, when we start from a magnetron orbit, ρ_{start}, much smaller than the trap size, the time required for the ions to diffuse out of the trap, τ_{trap} is practically independent of the trap size:

$$\tau_{trap} \approx \left[n\sigma \frac{\omega_z \omega_m}{\omega_c} \rho_{start} \right]^{-1} = \frac{\omega_c}{\omega_m} \tau_{start} \qquad \text{(II.C.1.9)}$$

where τ_{start} is the mean time between collisions right after the ions are loaded, $nsignma\,\omega_z \rho_{start}^{-1}$. For example, using this formula, N_2^+ ions in the presence of a background gas at a density of 3×10^4 cm^{-3} (corresponding to a room temperature background pressure of 10^{-12} T) and $\sigma \sim 10^{-12}$cm^2,[2] created 0.2 mm from the center of the trap, will remain trapped for about one month. (Note that if, instead of equipartition, we assume that the cyclotron orbit remains fixed—set presumably by the initial room temperature energy of the parent neutral—we get a much longer trapping time:

$$\tau_{trap} \approx \left[8 \frac{E_{room}}{E_{trap}} \frac{\omega_m}{\omega_c} n\sigma v_{room} \right]^{-1} \qquad \text{(II.C.1.10)}$$

which, for identical conditions as above yields a life time of two years!)

Thus the radial diffusion becomes a minor problem in the ultra-high vacuum regime in which we operate the trap [FLA87], especially when, as we shall see in Chapter VI, we have the additional ability to decrease the magnetron orbit at will.

Collisional Line Broadening

2. Obviously difficult to determine in general, this estimate of the cross-section seems sufficiently pessimistic. [PRI86]

Line broadening, on the other hand, can be a far more important problem, especially for precision cyclotron frequency measurement. In the worst case, the collision will entirely randomize the phase of the cyclotron motion and thus, in analogy to atomic line broadening (e.g., [COR77]), will cause of a broadening of the cyclotron resonance by $\approx \frac{1}{\tau_{collide}}$. There are two ways we can estimate the effects on a cyclotron measurement. We either use this line width to estimate the precision of a cyclotron measurement, or, instead, require that *no* collision take place during a cyclotron measurement.

To estimate how a broadened line will affect the precision of a measurement, we can use the standard rule of thumb that, when determining the central frequency, at best one may split a resonance line by its signal-to-noise ratio. In later chapters, I will discuss both the noise of the axial detector and how we use that detector to measure the cyclotron motion. Specifically, in Section VI.A, I will show that cyclotron motion with energy E_{cyc} can be transferred parametrically into axial motion with energy $E_z = \frac{\omega_z}{\omega_c} E_{cyc}$. In addition, in Section IV.A, I will present measurements of our present detector's noise. Using these facts, we can estimate the signal-to-noise ratio for a cyclotron measurement.

One way we express axial signals and noise is in units of the current that a single ion, moving endcap to endcap would produce. From Equation II.A.3.4, we know this current is $i = \frac{1}{2} e B_1 \omega_z$, currently $6.4 \times 10^{-14} A$. For historical reasons[3] we

call these "β"-units. For example $\beta = 0.1$ corresponds to the current a single ion would induce moving, at its peak, 10% of the trap size. We can also express the detector noise in these units, β_{noise}, which can then be compared quite easily to motions of the ion. (β_{noise}, then, has units of $Hz^{-1/2}$) In terms of these units, an experiment with integration time of $T_{\exp}$ will yield, on average, a noise equivalent to an rms ion excursion of $\beta_{noise}/(2T_{\exp}^{1/2})$. In the collisionally-broadened limit, then, the precision, P, of the experiment can be written:

$$P \equiv \frac{\Delta\omega_c}{\omega_c} \approx \frac{1}{\tau_{collide}} \frac{\beta_{noise}}{2T_{\exp}^{1/2}\beta_{signal}} \frac{1}{\omega_c} \qquad \text{(II.C.1.11)}$$

where the middle part of the expression on the right is the estimated signal-to-noise. Because $\tau_{collide}$ goes like β_{signal}^{-1}—the larger ion velocity makes collisions occur more frequently—we find that the critical background density, n, is independent of the size of the ion signal:

$$n \leq \left[\frac{2P}{\sigma\beta_{noise} z_0}\right]\left[\frac{\omega_c T_{\exp}}{\omega_z}\right]^{\frac{1}{2}} \qquad \text{(II.C.1.12)}$$

For a 10 second experiment, which will achieve a precision of 10^{-10} on an N_2^+ ion, using the present detector ($\beta_{noise} \approx 0.03\ Hz^{-1/2}$) and $\sigma \sim 10^{-12}\,\text{cm}^2$, we must have $n < 2\times10^5\ \text{cm}^{-3}$, corresponding to a room temperature pressure of $\approx 6\times10^{-12}$ T.

---

3. That is, for no reason anyone can remember.

On the other hand, we might require that *no* collisions occur during one cyclotron measurement. In that case, the effective width of the line will go like $T_{\exp}^{-1}$, and thus the precision will go like $T_{\exp}^{-3/2}$. Specifying a desired precision will then set the length of time require to achieve that precision:

$$T_{\exp} = \left[\frac{\beta_{noise}}{2\beta_{signal} P \omega_c} \right]^{2/3} \qquad \text{(II.C.1.13)}$$

We then require, on average, no collisions during that period; that is:

$$n < \frac{1}{\sigma v T_{\exp}} \qquad \text{(II.C.1.14)}$$

The velocity v will be set by β_{signal}, using the parametric detection scheme:

$$v = \beta_{signal} z_0 \sqrt{\omega_c \omega_z} \qquad \text{(II.C.1.15)}$$

In this case, for the critical density we get the rather ungainly expression:

$$n \leq \left[\frac{2P}{\beta_{noise}} \right]^{2/3} \frac{1}{\sigma z_0} \left[\frac{2}{\omega_z \omega_m} \right]^{1/6} \frac{1}{\beta_{signal}^{1/3}} \qquad \text{(II.C.1.16)}$$

This weak dependence on β_{signal} comes about because increasing the ion velocity decreases the mean time between collision. For $\beta_{signal} = 0.2$, with the same parameters as above, we will require $T_{\exp} = 10\ s$ and thus $n \leq 2\times10^5\ \text{cm}^{-3}$, too.

II.C.2 Impurity Ions

While neutral atoms can be a nuisance and may ultimately limit the attainable precision, background *ions* can be far more crippling, causing large frequency shifts, broadenings, and, it appears, temporal instabilities in all the trap motions. Using a

simple model, we can estimate the impact of these "bad" ions on the motion of our "good" ions.

For simplicity, let me introduce a "toy" model for the evolution of the z motion of two different ions, of mass m_1 and m_2, coupled by their electrostatic repulsion. In this model, I will assume that their radial separation is fixed at R, and that this separation is much larger than the extent of the z motion. These assumptions let us linearize the coulomb force in the z direction:

$$F_z = \frac{e^2}{R^3}(z_1 - z_2) \qquad \text{(II.C.2.1)}$$

where z_1 and z_2 are the positions of the two ions. Adding this force to the usual trapping force, we can write down the coupled equations of motion:

$$\ddot{z}_1 + \frac{eV_t}{m_1 d^2} z_1 = \frac{e^2}{m_1 R^3}(z_1 - z_2) \qquad \text{(II.C.2.2)}$$

$$\ddot{z}_2 + \frac{eV_t}{m_2 d^2} z_2 = \frac{e^2}{m_2 R^3}(z_2 - z_1)$$

and compute the normal modes and frequencies.

For $m_1 = m_2$, the coulomb interaction cancels for $z_1 + z_2$, as alluded to in Section II.A.3. (Our detector is only sensitive to $z_1 + z_2$.) However, when $m_1 \neq m_2$, we must solve the matrix equation:

$$\begin{bmatrix} \ddot{z}_1 \\ \ddot{z}_2 \end{bmatrix} = -\Omega^2 \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \qquad \text{(II.C.2.3)}$$

where:

$$\Omega^2 = \begin{bmatrix} \omega_1^2 - \dfrac{e^2}{m_1 R^3} & \dfrac{e^2}{m_1 R^3} \\ \dfrac{e^2}{m_2 R^3} & \omega_2^2 - \dfrac{e^2}{m_2 R^3} \end{bmatrix} \tag{II.C.2.4}$$

and

$$\omega_i = \frac{eV_t}{m_i d^2} \tag{II.C.2.5}$$

The usual way to solve this matrix equation is to find a new set of coordinate, z_A and z_B, related to the original coordinates by a unitary transformation, U:

$$\begin{bmatrix} z_A \\ z_B \end{bmatrix} = U \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \tag{II.C.2.6}$$

such that, in the new coordinates, Ω^2 is diagonal; that is, the coordinates z_A and z_B are uncoupled. Plugging (II.C.2.6) back into (II.C.2.3), we see that U must satisfy:

$$\begin{bmatrix} \omega_A^2 & 0 \\ 0 & \omega_B^2 \end{bmatrix} = U\,\Omega^2\,U^{-1} \tag{II.C.2.7}$$

where ω_A and ω_B are the frequencies of the uncoupled modes. Thus the eigenvalues of Ω^2 are the squared normal mode frequencies and the eigenvectors of Ω^2 form the rows of U.

In the simplest case, the frequency difference between the masses is much larger than the repulsion; that is, $|\omega_1^2 - \omega_2^2| \gg \dfrac{e^2}{m_{1,2} R^3}$. In this case, to lowest order, we find:

$$\omega_A \approx \omega_1 - \frac{e^2}{2m_1R^3\omega_1} \qquad \text{(II.C.2.8)}$$

$$\omega_B \approx \omega_2 - \frac{e^2}{2m_2R^3\omega_2}$$

Quite remarkably, the resulting shift in resonance frequency is independent of the perturbing ion, and always toward lower frequency. (This result is reasonable: the coulomb interaction is repulsive and always has the effect of weakening the spring constant of the trap.) For a single perturbing ion, with a magnetron orbit 1 mm away, for N_2^+ we expect from (II.C.2.8) a shift in ν_z of 0.4 Hz—about eight times its width!

However, when the ions are close together, or have nearly the same mass, the coulomb force can have an even larger effect and cause the ions' motion to "lock" together. In that case, we must use (II.C.2.4) to write down the eigenfrequencies, in general:

$$\omega^2 = \frac{1}{2}\left\{\frac{K_0 - K_1}{\bar{m}} \pm \sqrt{\frac{4K_1^2}{m_1m_2} + \left[\frac{(K_0 - K_1)\Delta m}{m_1m_2}\right]^2}\right\} \qquad \text{(II.C.2.9)}$$

where:

$$K_0 \equiv \frac{eV_t}{d^2} = \text{spring constant of the trap}$$

$$K_1 \equiv \frac{e^2}{R^3} = \text{spring constant of the coulomb repulsion}$$

$$\Delta m \equiv m_2 - m_1$$

and $$\frac{1}{\overline{m}} \equiv \frac{1}{2}\left[\frac{1}{m_1} + \frac{1}{m_2}\right]$$

When the first term under the radical in (II.C.2.9) dominates, the axial modes will be closer to the symmetric/anti-symmetric combinations $\frac{1}{\sqrt{2}}(z_1 \pm z_2)$ than the separate ion modes, z_1 and z_2. That is, for $K_1 \gg \frac{1}{2}K_0\frac{\Delta m}{\overline{m}}$, the coulomb repulsion will cause the ions' motion to lock together into (correlated) symmetric and anti-symmetric motions. Although the anti-symmetric mode always has the lower frequency, the total current induced in our detector is proportional to $z_1 + z_2$, and thus we can observe only the symmetric mode. In the strongly-coupled case, that mode will have the frequency:

$$\omega^2 \approx \frac{K_0}{\overline{m}} - \frac{1}{8}\frac{K_1}{\overline{m}}\frac{\Delta m^2}{\overline{m}^2} \qquad \text{(II.C.2.10)}$$

The left-over effect of the coulomb interaction on the frequency of this mode is almost non-existent: the coulomb coupling K_1 appears reduced by the factor $\frac{1}{8}\frac{\Delta m^2}{\overline{m}^2}$. Since this modes is the only one we detect, we will see the two ions as if they had a mass equal to the harmonic mean of the two locked masses, $\overline{m}$. Equation II.C.2.10 also shows that identical ions ($\Delta m = 0$, $\overline{m} = m_1 = m_2$) show no shift in frequency. This

last fact is true more generally. For any cloud consisting of only one ion species, whose ion-ion interactions can be described by a central potential, the center-of-mass mode will have the same frequency as the single, isolated ion.

In several different instances, we expect coulomb coupling will have a bad effect on the axial mode of ions we wish to detect ("good ions"[4]). In the first case, other species in the trap ("bad ions" made simultaneously with the good ions) can cause temporally varying frequency shifts. Right after their creation, these bad ions probably will have large axial orbits (because they will not be cooled effectively by the detector), will be weakly coupled to the good ions, and hence will shift the trap's spring constant by $K_0 - K_1$. As these bad ions slowly cool, (by collision, we assume, since they are far away from resonance and thus can not be cooled directly by the detector) the coulomb force will increase, causing a small, time-dependent shift of the good ions' resonance. Eventually, when they cool sufficiently, the good and bad ions may couple together and oscillate jointly far away from the expected resonance frequency of the isolated good ions.

In the second case, we might have only one species in the trap, but have several ions with very large magnetron (or even cyclotron) orbits. ("bad orbits") Because of anharmonicity (Section II.B.3), we know that these orbits may have a slightly different

4. Value judgements like these are inescapable, and such labels usually originated late at night

frequency than ions in the center of the trap. We might consider this frequency shift, $\delta\omega$, as if it were due to a slight mass shift, $\delta m = \frac{2m}{\omega}\delta\omega$. For example, for $C_4 = 10^{-3}$, an N_2^+ ion in a magnetron orbit of 1 mm will have an effective $\frac{\delta m}{m} = 10^{-4}$. In this case, too, we might get both frequency shifts and frequency locking; that is, even though the good ion may be driven harmonically, because of the coupling, its frequency will be pulled toward the ion in the bad orbit. That is, the immunity to coulomb shifts that the (measured) resonant frequency of a cloud of identical ions usually possesses is destroyed by anharmonicity.

A final comment about ion-ion coupling. We once envisioned comparing *simultaneously trapped*, nearly degenerate ions (like ^3He and ^3H, for which $\frac{\delta m}{m} = 6\times10^{-6}$) The calculation above indicates that we would have to keep them separated by $R \gg 2.5$mm in order to avoid frequency locking. Since the trap's radius is only 7 mm, we would have to move one ion to the center and other right to the radial fringe of the trap! More likely, then, when we try that experiment, we will compare *each* ion to *HD* ($\frac{\delta m}{m} = 2\times10^{-3}$), for which we need only $R \gg 0.35$ mm, a much more practical separation.

II.C.3 Patch Effects

One possible source of field imperfections is small patches of charge lying on the trap electrodes. These patches could disrupt both reflection and azimuthal

symmetry and could cause all sorts of shifts due to higher order anharmonicity. Because our electrodes are gold-plated, we must expect some patches to occur [REF], and have taken steps to minimize these patches.

For example, a small patch of charge Q located on the upper endcap near the hole through which the atoms enter the trap will generate a potential, near the center of the trap:

$$\Phi_{patch} = \frac{Q}{\sqrt{\rho^2 + (z - z_0)^2}} \qquad \text{(II.C.3.1)}$$

We can expand this potential in a Legendre series (trivially, it turns out, see [JAC75], p 93.):

$$\Phi_{patch} = \frac{Q}{z_0} \sum_l \left[\frac{r}{z_0}\right]^l P_l(\cos\theta) \qquad \text{(II.C.3.2)}$$

where r and θ are the usual spherical coordinates. Thus a patch could effect the trap dynamics at all orders.

For example, the $l = 1$ term is equivalent to an additional voltage imbalance, ΔV_A, between the upper and lower endcaps:

$$\Delta V_A = \frac{2Q}{B_1 z_0} \qquad \text{(II.C.3.3)}$$

which will shift the equilibrium position of ion. The $l = 2$ term effectively changes the resonant frequency as if an additional trapping voltage ΔV_T, were added:

$$\Delta V_T = \frac{2Q}{z_0}\frac{d^2}{z_0^2} \qquad \text{(II.C.3.4)}$$

and so on.

We could detect the presence of a patch several different ways. We can measure the trapping potential required to bring different species into resonance with our detector. In the absence of a patch, this voltage will be directly proportional to the mass. Therefore, when we plot the trapping potential against the mass, if we find a non-vanishing y-intercept, then we must have a patch causing a ΔV_T. Alternatively, we can measure the product $B_1 B_3$ using the technique summarized by Equation II.B.14, plotting the axial frequency shift against the applied endcap voltage. The resulting parabola has a curvature which gives $B_1 B_3$ and a minimum, in the absence of a patch, at $V_- = 0$. If the minimum of the parabola is not at zero, we may in fact be compensating for a non-zero patch effect, although such a shift could also be caused by an asymmetry in the spacing of the two endcaps. However, this mechanical asymmetry should cause a shift that depends on the mass of the trapped ions (because the shift should a constant fraction of the trapping potential) whereas a patch-induced asymmetry should cause a shift which is roughly independent of the mass.

Although the $l = 1$ and $l = 2$ terms can be canceled out by additional potentials on the lower endcap and the ring, the higher-order anharmonic terms will remain and, as we've discussed in Section II.B.2 and II.B.3, we would prefer to eliminate such terms. In addition, any time variation in the patches would be another nuisance.

Therefore, as I will discuss in Chapters V and IV, we have taken steps to eliminate charged patches.

II.C.4 Anharmonicity Revisited

While we have already discussed the effect of axial anharmonicity on the axial motion, we have neglected the effects of non-harmonic electric fields on the other trap modes. Using the discussion of anharmonicity in Section II.B.3 as a starting point, we can see that C_4-terms can cause shifts in all the modes and these shifts, in turn, depend on the amplitudes of all the modes.

As discussed in Section II.B.3, the dominant imperfection in the trapping potential is the quartic term:

$$\Phi_s = \frac{1}{2} C_4 \frac{z^4 - 3z^2\rho^2 - 3\rho^4/8}{d^4} \qquad \text{(II.C.4.1)}$$

In the presence of this perturbation, the effective axial frequency becomes:

$$\omega_z^2 \rightarrow \omega_z^2 \left[1 - 3C_4\frac{\rho^2}{d^2} + 2C_4\frac{z^2}{d^2}\right] \qquad \text{(II.C.4.2)}$$

As described in that section (Equation II.B.18), the z^2 term causes a shift in the axial frequency due to the non-linearity:

$$\frac{\Delta\omega_z}{\omega_z} \approx \frac{3}{4} C_4 \frac{z^2}{d^2} \qquad \text{(II.C.4.3)}$$

In addition, the middle term in (II.C.4.2) gives a radius-dependent shift to the axial frequency:

$$\frac{\Delta\omega_z}{\omega_z} \approx -\frac{3}{2} C_4 \frac{\rho_m{}^2 + \rho_c{}^2}{d^2} \qquad \text{(II.C.4.4)}$$

and thus the total axial frequency shift will be:

$$\frac{\Delta\omega_z}{\omega_z} \approx \frac{3}{4} \frac{C_4}{d^2} \left[z^2 - 2\rho_m{}^2 - 2\rho_c{}^2 \right] \qquad \text{(II.C.4.5)}$$

Thus, for $C_4 = 10^{-4}$, we expect a 1 mm magnetron orbit will cause about a 1 Hz shift in the axial frequency of an N_2^+ ion. In fact, this residual anharmonicity has been used to detect changes in magnetron orbit size [MOO87].

The potential Φ_s also will shift the cyclotron and magnetron frequencies. To compute these shifts, we can calculate the force due to this additional radial electric field using (II.C.4.1):

$$\vec{F}_\rho = 3m\omega_z{}^2 C_4 \left[\frac{z^2 - \rho^2/4}{d^2} \right] \vec{\rho} \qquad \text{(II.C.4.6)}$$

For an oscillating axial motion, $\langle z^2 \rangle = \frac{1}{2} z_{peak}^2$, and thus we can re-write (II.A.2.6) to include this perturbation:

$$\ddot{\vec{\rho}} - \omega_c \hat{z} \times \dot{\vec{\rho}} - \frac{1}{2}\omega_z{}^2 \left[1 + 3C_4 \frac{z_{peak}^2 - \rho^2/2}{d^2} \right] \vec{\rho} = 0 \qquad \text{(II.C.4.7)}$$

For positive C_4, the z_{peak} term effectively increases ω_z, and, using (II.A.2.13), we find:

$$\frac{\Delta\omega_m}{\omega_m} = 3C_4 \frac{z_{peak}^2}{d^2} \qquad \text{(II.C.4.8)}$$

and

$$\frac{\Delta\omega_c}{\omega_c} = -3C_4\frac{\omega_m}{\omega_c}\frac{z_{peak}^2}{d^2} \qquad \text{(II.C.4.9)}$$

The cyclotron shifts are smaller by a factor of $\frac{\omega_m}{\omega_c}$ than either the axial or magnetron shifts, as we might expect: the cyclotron motion is mostly magnetic; the electric fields already are a small perturbation.

The $\rho^2\vec{p}$ term in (II.C.4.7) causes radial anharmonic shifts, just like the z^3 shifts in section II.B.3. These shifts due to the radial motion are a bit more involved, especially the effect of the cyclotron orbit on the magnetron motion (and vice versa), and the calculation is not particularly enlightening. I will therefore use the results from Brown and Gabrielse [BRG86], which, including the shifts we have discussed above, summarize the effects of C_4 anharmonicity on the trap modes:

$$\frac{\Delta\omega_c}{\omega_c} \approx \frac{3C_4}{d^2}\frac{\omega_m}{\omega_c}\left[\frac{1}{2}\rho_c{}^2 - z^2 - \rho_m{}^2\right] \qquad \text{(II.C.4.10)}$$

$$\frac{\Delta\omega_z}{\omega_z} \approx \frac{3C_4}{4d^2}\left[-2\rho_c{}^2 + z^2 - 2\rho_m{}^2\right]$$

$$\frac{\Delta\omega_m}{\omega_m} \approx \frac{3C_4}{d^2}\left[-\rho_c{}^2 + z^2 - \frac{1}{2}\rho_m{}^2\right]$$

Thus a small anharmonicity in the trapping potential is not likely to have a dramatic effect on the cyclotron measurement.

II.C.5 Magnetic Inhomogeneity

Up to this point, I've mentioned only electric field inhomogeneities. In addition to these, there can be magnetic field inhomogeneity both in the magnetic field of our original magnet and, more importantly, because of fields produced by the susceptibility of the trap materials. While the magnetic field (without the trap in place) had a measured homogeneity of about 10^{-8} over a 1 cm^3 volume, we have estimated [FLA87] that the 85000 gauss field should induce, at trap center, a second order inhomogeneity of about 6 gauss/cm^2. Thus the trap materials produce a "magnetic bottle" four orders of magnitude larger than the inhomogeneity without the trap. While bottles have been used extensively at Washington to detect the cyclotron motion of electrons [WED73,BRG86], we consider this bottle to be a problem and, eventually, will take steps to eliminate it.

We can calculate fairly directly the effect on the trap frequencies that a magnetic bottle induces. The field, $\Delta\vec{B}$, associated with a dipole bottle of strength B_2 is:

$$\Delta\vec{B} = B_2\left[(z^2 - \frac{\rho^2}{2})\hat{z} - z\rho\hat{\rho}\right] \qquad \text{(II.C.5.1)}$$

We can compute its effect on the axial motion by calculating the extra axial force, ΔF_z, that the bottle induces:

$$\begin{aligned}\Delta F_z &= \frac{e}{c}(\vec{v}\times\Delta\vec{B})\cdot\hat{z} \qquad \text{(II.C.5.2)}\\ &= -\frac{e}{c}zB_2\hat{\rho}\cdot(\hat{z}\times\vec{v})\end{aligned}$$

Since the magnetron and cyclotron orbits are circular motions, we can use $\hat{z}\times\vec{v} = \omega_{m,c}\vec{\rho}_{m,c}$ and write:

$$\Delta F_z = -\frac{e}{c} B_2 \left[\omega_m \rho_m{}^2 + \omega_c \rho_c{}^2 \right] z \qquad \text{(II.C.5.3)}$$

Thus the bottle can cause an additional harmonic force and, for $B_2 > 0$, strengthens the trapping potential. Assuming this force is a small perturbation on the axial motion, the bottle will induce a frequency shift:

$$\frac{\Delta\omega_z}{\omega_z} \approx \frac{1}{4} \frac{B_2}{B} \left[\rho_m{}^2 + \frac{\omega_c}{\omega_m} \rho_c{}^2 \right] \qquad \text{(II.C.5.4)}$$

While the shift due to the magnetron is quite small (a 1 mm orbit for an N_2^+ ion gives only an 0.03 Hz shift which will likely be overshadowed by shifts to an anharmonic *electric* field), the effect of the cyclotron motion can be quite large. (Similar shifts have been used to measure cyclotron resonances of electron. [WED73]) As an example of the size of this shift, in our parametric pulse schemes, a starting z orbit of 1 mm will cause a 1.6 Hz shift in the axial resonance frequency when the z motion has been parametrically driven into the cyclotron mode.

This magnetic inhomogeneity will also shift the magnetron and cyclotron resonances, and we could calculate those shifts using the same procedure we used for the axial shifts. Rather than going through that algebra, I will just quote the results from a similar calculation [BRG86]:

$$\frac{\Delta\omega_c}{\omega_c} \approx \frac{1}{2} \frac{B_2}{B} \left[z^2 - \rho_m{}^2 - \rho_c{}^2 \right] \qquad \text{(II.C.5.5)}$$

and

$$\frac{\Delta\omega_m}{\omega_m} \approx -\frac{1}{2}\frac{B_2}{B}\left[z^2 - \rho_m{}^2 - \frac{\omega_c}{\omega_m}\rho_c{}^2\right] \qquad \text{(II.C.5.6)}$$

A 1 mm z orbit could cause a 3×10^{-7} shift in the cyclotron frequency. This shift, though small, is still several orders of magnitude greater than the ultimate precision we seek. In addition, in analogy to the anharmonic axial resonances of Section II.B.3, and the anharmonic cyclotron effect mentioned in the previous section, this frequency shift with $\rho_m{}^2$ (and z^2, for that matter) can cause hysteresis and sweep-direction dependence in swept cyclotron resonances.

II.C.6 Magnetic Field Drifts

The central tenet of our experiment is that we can compare masses by comparing the cyclotron frequencies of ions in the same magnetic field. Because it is difficult to measure simultaneously the cyclotron frequencies of two different species, short term jitter and longer term drift of the magnetic field will translate directly into uncertainties in the mass ratio. Therefore, the stability of the magnet is very important.

Our particular magnet was selected by the manufacturer to have exceptionally small drift [FLA87], and our preliminary measurements indicate a long-term stability of better than 10^{-9} per hour. This decay should be linear, and, assuming our magnet performs at least as well as earlier magnets, the uncertainty in this decay should be less than 10^{-11} per hour. [VMF85]

However, short term fluctuations, due for example to moving elevators, changing atmospheric pressure, etc., probably also contribute around the 10^{-9} level. These

fluctuations, too, will mimic fluctuations in the mass. There is no simple prescription to eliminate these fluctuations, though the superconducting solenoid itself should have some shielding effect against external fields. We may supplement this shielding using additional superconducting shields [GAB88] (which try to keep the flux threading them constant) or try to develop suitable technology to measure these fluctuations directly.

II.C.7 Trap Asymmetry and Tilts

In this section, I will discuss the effect of azimuthal asymmetry in the trap electrodes and misalignment between the magnetic and electrostatic axes. These last two perturbations represent both the largest and, as it turns out, the least important perturbations in the trap—least important because we can give a prescription that compensates for these effects to all orders. This prescription enables us convert the three measured trap frequencies of the ions into a robust, free-space cyclotron value, proportional to its $(\text{mass})^{-1}$.

We first discuss the asymmetry, and then add the magnetic field tilt into the same formalism. A simple counting argument shows that we need only one parameter to specify the quadratic, azimuthal asymmetry in the electrodes. The quadratic potential for the trap must be given by a symmetric, traceless, 3×3 tensor. (Symmetric because the anti-symmetric part has no physical effect, and traceless because, in a source-free region, $\vec{\nabla}^2\Phi = 0$.) Thus, in general, there are $3 \cdot 2 \cdot 1 - 1 = 5$ possible parameters in the potential. The choice of axes is arbitrary, and, for convenience, we

can rotate to the principal axes, for which the tensor is diagonal. Thus, of the remaining 5 parameters, 3 are arbitrary (the three Euler angles of the rotation that gets us to the principal axes). Of the two remaining parameters, one gives the overall strength of the potential, leaving us with one remaining parameter. Following the description by Brown and Gabrielse, (who first published the technique [BRG82]) we call the left over paramenter ε, and let it represent an ellipsoidal asymmetry:

$$\Phi = \frac{m}{2e}\omega_z{}^2\left[z^2 - \frac{1}{2}(x^2+y^2) - \frac{1}{2}\varepsilon(x^2-y^2)\right] \qquad \text{(II.C.7.1)}$$

When written in cylindrical coordinates, we see explicitly that ε breaks the azimuthal symmetry:

$$\Phi = \frac{m}{2e}\omega_z{}^2\left[z^2 - \frac{1}{2}\rho^2(1-\varepsilon\cos 2\phi)\right] \qquad \text{(II.C.7.2)}$$

In this formalism, we also can treat a tilt of the magnetic field with respect to the electrostatic principal axes. Since we have specified these axes, the magnetic tilt has two angles: θ and ϕ, the usual spherical angles. (The angle ϕ becomes important because it couples to the asymmetry ε.) Using rectangular coordinates for the moment, we can write down the ions' equations of motion, taking ε, θ and ϕ explicitly into account:

$$\begin{bmatrix}\ddot{x}\\ \ddot{y}\\ \ddot{z}\end{bmatrix} - \omega_c\begin{bmatrix}\dot{y}\cos\theta - \dot{z}\sin\theta\sin\phi\\ \dot{z}\sin\theta\cos\phi - \dot{x}\cos\theta\\ \dot{x}\sin\theta\sin\phi - \dot{y}\sin\theta\cos\phi\end{bmatrix} + \omega_z{}^2\begin{bmatrix}-\frac{1}{2}(1+\varepsilon)x\\ -\frac{1}{2}(1-\varepsilon)y\\ z\end{bmatrix} = 0 \qquad \text{(II.C.7.3)}$$

and solve for the normal modes. Because Brown and Gabrielse explicitly do this, I will only present their results.

We can label the three normal modes of the tilted, asymmetric trap $\overline{\omega}_z$, $\overline{\omega}_m$ and $\overline{\omega}_c$. We get these modes by trying solutions like $e^{-i\omega t}$, and solving the determinant equations that results from (II.C.7.3). The most important relation between these frequencies is:

$$\omega_c^2 = \overline{\omega}_c^2 + \overline{\omega}_z^2 + \overline{\omega}_m^2 \qquad \text{(II.C.7.4)}$$

Quite remarkably, even in the presence of tilts and asymmetry, the quadratic sum of the normal modes, to all orders in ε, θ and ϕ, remains ω_c^2. Thus, this prescription provides a robust, high-precision value for the true cyclotron frequency.

We also can estimate the relationship between the measured trap modes and the ideal axial and magnetron modes:

$$\overline{\omega}_z \approx \omega_z \left[1 - \frac{3}{4}\sin^2\theta\left(1+\frac{1}{3}\varepsilon\cos 2\phi\right)\right] \qquad \text{(II.C.7.5)}$$

and

$$\overline{\omega}_m \approx \frac{1}{2}\frac{\overline{\omega}_z^2}{\overline{\omega}_c}\left(1-\frac{1}{2}\varepsilon^2\right)\left(1+\frac{9}{4}\sin^2\theta\right) \qquad \text{(II.C.7.6)}$$

Since there is no way to measure the a priori ("untilted") axial frequency, Equation II.C.7.5 cannot be used to measure the tilt. On the other hand, the magnetron frequency can be used for this measurement. Assuming that the asymmetry is less important than the tilt (a reasonable assumption [BRG82] which seems to hold for

most measured traps, unless the trap has been dropped on its side), we can use the three measured modes to give the tilt between the trap's z axis and the magnetic field:

$$\theta = \sin^{-1}\left[\frac{2}{3}\sqrt{\frac{2\bar{\omega}_m\bar{\omega}_c}{\bar{\omega}_z^{\,2}} - 1}\right] \qquad \text{(II.C.7.7)}$$

For example, a 1° tilt will cause (an easily measured) ≈ 2 Hz shift between $\bar{\omega}_m$ and $\tilde{\omega}_m \equiv \frac{1}{2}\frac{\bar{\omega}_z^{\,2}}{\bar{\omega}_c}$ for N_2^+ ions, a shift of about 10^{-3}. (In this sense, the tilts are, by far, the largest perturbation.)

Brown and Gabrielse also give a convenient power series expression for ω_c in terms of the measured modes:

$$\frac{\omega_c}{\bar{\omega}_c} = 1 + \frac{1}{2}\left[\frac{\bar{\omega}_z}{\bar{\omega}_c}\right]^2 + \frac{1}{8}\left[\frac{\bar{\omega}_z}{\bar{\omega}_c}\right]^4\left[\left[\frac{\bar{\omega}_m}{\tilde{\omega}_m}\right]^2 - 1\right] + \cdots \qquad \text{(II.C.7.8)}$$

where $\tilde{\omega}_m$ is the "if the trap weren't tilted" magnetron frequency given above.

We see from this expression that, for measurements of a certain precision, $\bar{\omega}_m$ may not have to be measured at all, depending, of course, on the degree of tilt and asymmetry. In terms of those parameters (and assuming that they are relatively small):

$$\frac{\omega_c}{\bar{\omega}_c} = 1 + \frac{1}{2}\left[\frac{\bar{\omega}_z}{\bar{\omega}_c}\right]^2 + \frac{9}{16}\left[\frac{\bar{\omega}_z}{\bar{\omega}_c}\right]^4\left[\sin^2\theta - \frac{2}{9}\varepsilon^2\right] + \cdots \qquad \text{(II.C.7.9)}$$

For N_2^+ ions ($\frac{\bar{\omega}_z}{\bar{\omega}_c} \approx 3.5\times10^{-2}$), a 1° tilt will cause less than a 3×10^{-10} error if we ignore the magnetron measurement entirely. For $^3He^+$ ions, ($\frac{\bar{\omega}_z}{\bar{\omega}_c} \approx 3.7\times10^{-3}$), a tilt

as large as 3° would cause less than a 3×10^{-13} error! Thus, for high precision measurements on lighter ions, the magnetron frequency will only be important to cool the ions. (See Chapter VI) An exact measurement of its frequency, however, will not be important when calculating the free-space cyclotron frequency.

II.C.8 Relativistic Shifts

Special Relativity provides a final "perturbation" to our measurements, or rather, a perturbation to our theory. All of our derivations in Section II.A.2 assumed that we could apply classical mechanics completely and exclusively. However, for typical single ion measurements, we expect the cyclotron orbits of lighter ions (like $^3He^+$ ions) to have v/c as large as 3×10^{-5}, and thus relativistic corrections will be important at the 10^{-9} level.

The relativistic shifts in the cyclotron frequency are fairly simple to calculate. Since the cyclotron motion is perpendicular to the magnetic field, we expect relativistic shifts due only to the increase of the ion's effective mass, $m \rightarrow \left[1-v^2/c^2\right]^{-1/2} m$. The cyclotron frequency depends inversely on the mass, (II.A.2.7); therefore, we expect $\frac{\Delta\omega_c}{\omega_c} = -\frac{\Delta m}{m}$. Averaging over the axial orbit and neglecting shifts due to the magnetron motion[5] we have:

$$\frac{\Delta\omega_c}{\omega_c} = -\frac{1}{2}\frac{{\omega_c}^2}{c^2}\left[{\rho_c}^2 + \frac{\omega_m}{\omega_c}z^2\right] \qquad \text{(II.C.8.1)}$$

Even for strongly excited N_2^+ ions, using detection schemes outlined in the next chapter, the shifts due to the cyclotron and axial motions should be about the same, roughly 3×10^{-12}. However, these shifts scale like m^{-2}, so the shift for $^3He^+$ will be closer to 3×10^{-10}. For very light particles (electrons), Gabrielse [GDK85] has observed hysteresis (as per Section II.B.3) due to this relativistic mass shift.

Determining the axial frequency shift is complicated by Lorentz contraction. Brown and Gabrielse, using the large-n limit of a relativistic quantum calculation show, for the axial shift:

$$\frac{\Delta\omega_z}{\omega_z} = -\frac{1}{4}\frac{{\omega_c}^2}{c^2}\left[{\rho_c}^2 + \frac{3}{8}\frac{\omega_m}{\omega_c}z^2\right] \qquad \text{(II.C.8.2)}$$

We see that the *relative* axial shift is of the same order as the cyclotron shift; therefore the absolute axial shift will always be negligible.

Summary of Perturbations

In these sub-sections, I have outlined the most important perturbations to the ideal Penning trap modes. We will return to these shifts after we discuss detection techniques (Chapters V and VI) when we will have better understanding of the relative

amplitudes in the different modes. In Chapter VII, we will use these calculated shifts to estimate the present systematics, (for example, the best C_4 and bottle field) and return conclude Chapter VII with an estimate of how well we ultimately can measure mass ratios in the presence of the remaining perturbations.

CHAPTER III

DETECTOR THEORY AND IMPLEMENTATION

This chaper addresses two crucial ideas central to our ion detection schemes: SQUIDs and FFTs. I have grouped them together because, though central, they are sufficiently unfamiliar in detail to most physicists to warrant an explanation. In addition, these two ideas are logically linked: both represent crucial steps between the trapped ions and the experimenters. The SQUID (and the associated cryogenic electronics) provides the physical amplification required to detect the velocity of the trapped ion. The FFT (and associated digital signal processing) provides the software to interpret these physical signals in some of the convenient ways discussed in Chapter III. Furthermore, both the SQUID and the FFT impose certain constraints on the system, especially the signal-to-noise, and understanding these constraints in the current system is crucial for future improvements. In both cases, then, I have tried to include enough theory and explanation so that the reader can understand the specific application to our experiment.

In another sense, too, this chapter serves as a bridge between the theory and the results. Each section must, by necessity, begin to deal with the specific implementations chosen for our experiment, and hence I begin to discuss some of the experimental "nuts and bolts" of ion detection. For example, I will discuss the high-Q tuned circuits we used with the SQUID because they are involved intimately in its optimization for ion detection. Also, I will explain the specific algorithm for variance

control required to make an effective spectrum analyzer using the FFT.

III.A RF SQUIDs

In this section, I will describe how a superconducting ring, closed by a "weak link" can be used as a very sensitive current amplifier. The earliest paper that outlines this use of the Josephson effect appeared in 1967 [SIZ67], and the device became a focus of attention for the next decade. By the end of that decade, the operating theory and a reasonably sophisticated treatment of the noise processes in an RF SQUID had been fairly well formulated and tested. (Many of the noise theory papers were written as part of an effort to use RF SQUIDs in gravitational wave detectors [HPE77,GUR77,JAB75].)

In the first part of this section, I will discuss how a Josephson Junction can be used as an audio-frequency flux detector. In the second part, I describe how we modified this flux detector to make it a sensitive enough to detect the current of a single, trapped ion. Most of this second section is a paper we published in the Journal of Applied Physics.

III.A.1 RF SQUIDs and Parametric Amplification

The principle effect of incorporating a weak link into a superconducting ring is to make the circulating supercurrent periodic in the flux threading the loop [FLS64,TIN75]):

$$J_s = -J_c \sin\frac{2\pi\Phi}{\Phi_0} \qquad \text{(III.A.1.1)}$$

where J_c is the critical current of the weak link, Φ is the flux threading the loop, and Φ_0 is the flux quantum, $h\frac{c}{2e} \simeq 2\times10^{-7}$ gauss-cm^2. (Another nice, quantum mechanical derivation of this result is in [FLS64].) The minus sign in III.A.1.1 indicates that the induced current opposes flux changes through the loop. The total flux threading the loop is the sum of the external flux and the flux set up by supercurrent circulating in the loop. Since the current, J_s in a ring of inductance L gives rise to a flux LJ_s, the total flux threading the loop is:

$$\Phi = \Phi_x - LJ_c \sin\frac{2\pi\Phi}{\Phi_0} \qquad \text{(III.A.1.2)}$$

This constant, LJ_c, determines the behavior of the loop. If it is very small ($\ll\Phi_0$), then $\Phi\simeq\Phi_x$, and the weak link effectively destroys flux quantization by destroying the coherence of the supercurrent. If, on the other hand, $LJ_c\to\infty$, then Φ becomes independent of Φ_x, and, it turns out, can only be a multiple of Φ_0. In this case, the weak link is so weak as to have no effect at all, and the ring behaves like a perfectly superconducting ring, opposing any applied flux by setting up the appropriate supercurrent. Figure III.A.1 shows a plot of Φ vs. Φ_x for $LJ_c = 2\pi\Phi_0$, a typical value for commercial sensors. Note that equation III.A.1.2 ignores the resistance and capacitance of the junction. This assumption is valid if the time scale of flux changes is $\ll \frac{R_s}{L_s}$, a constraint which is always satisfied in our system. [JAB75]

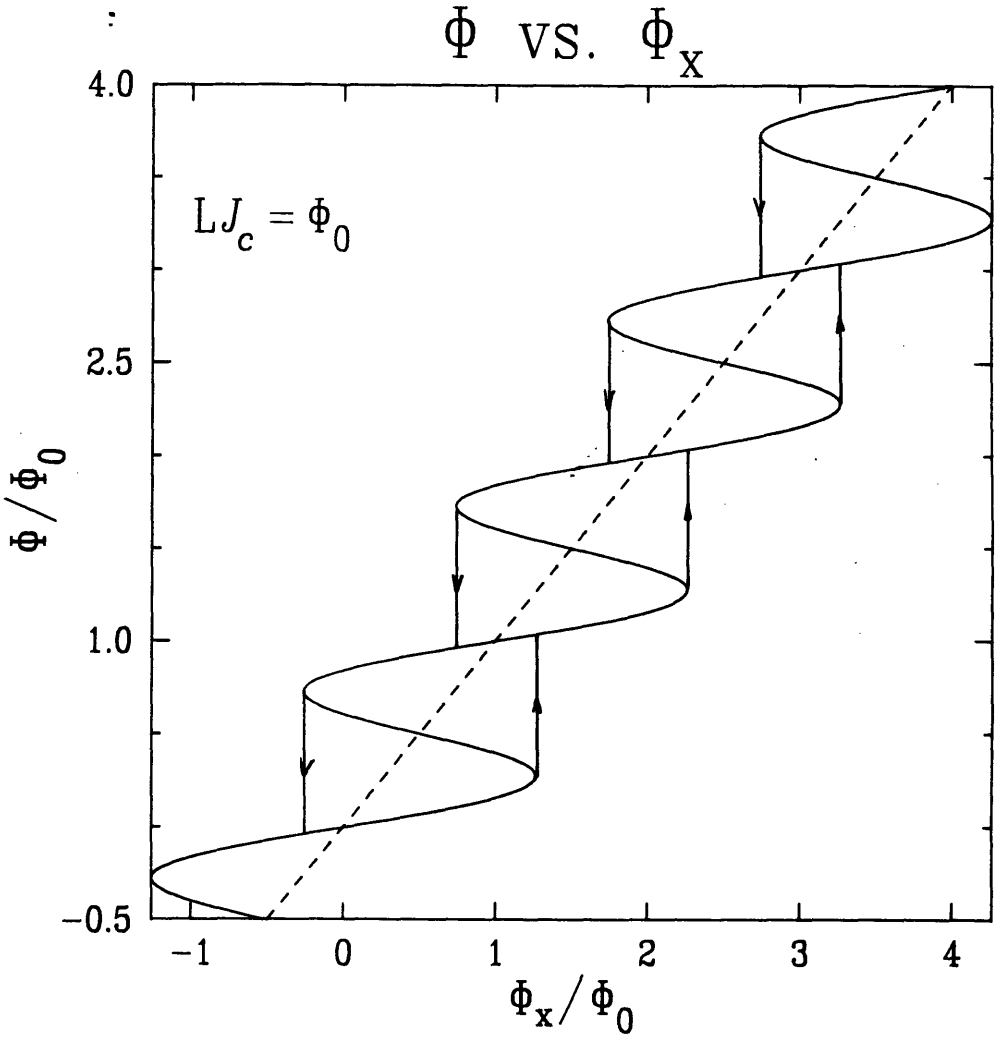

Figure III.A.1.1. Graph of Φ vs. Φ_x (Equation III.A.1.2). The parts of the curve with negative slope correspond to unstable conditions. Thus the flux follows the hysteretic paths indicated.

I will use the Φ vs. Φ_x graph to explain the dynamics of the RF SQUID. Kurkijaarvi [KUR72] has shown that the parts of the curve with negative slope correspond to unstable behavior. That is, as Φ_x is changed, Φ cannot remain on the parts of the curve with negative slope, and Φ will jump to the nearest, stable value. Thus, like the non-linear ion response discussed in Chapter II, this system exhibits a characteristic hysteresis.

We can use figure III.A.1.1 to determine the energy loss associated with one hysteresis loop. [TIN75] In general, the energy is the time integral of the power dissipated over one cycle:

$$\Delta W = \oint I\, V\, dt \qquad \text{(III.A.1.3)}$$

where I and V are the current and voltage around the loop. For an inductive device (like a SQUID loop), $V = \mathrm{L}\dot{I}$, so $V\,dt = \mathrm{L}\,dI$. In addition, as in equation III.A.1.2, $\mathrm{L}I = \Phi_x - \Phi$. Thus, using $dI = \frac{1}{\mathrm{L}} d\Phi$ and the fact that $\oint \Phi d\Phi = 0$ around any closed loop, we have:

$$\Delta W = \frac{1}{\mathrm{L}} \oint \Phi_x\, d\Phi \qquad \text{(III.A.1.4)}$$

Thus the area of a loop on Fig III.A.1.1 is proportional to the energy dissipated around the loop. (This energy is most likely dissipated as microwave energy through the weak link [LOU74]). For many practical devices, this energy is ~ Φ_0^2 / L, typically 4×10^{-21} j.

The standard way of exploiting this relationship for sensitive flux detection [SIZ67] is to couple an RF-tuned circuit loosely to the SQUID loop. (see Figure III.A.1.2) By measuring the voltage across the excitation tank circuit, the energy-loss discussed above can be detected. The peak voltage reaches a plateau, and the level of this plateau depends on the external flux, Φ_x. This plateau can be explained quite simply.

Assume that the constant external flux through the loop is such that there is no "static" supercurrent flowing; for example, whenever $\Phi_x = \Phi$. When we supply an RF current, I_{RF}, to the resonant tank circuit, (neglecting the presence of the SQUID loop for a moment), the current circulating in the tank builds up to a steady-state value $I_{tank} = QI_{RF}$, where Q is the usual quality factor of the tank circuit. This current, in turn, induces a current in SQUID loop, $I_{SQ} = \frac{M}{L_{SQ}} I_{tank}$. Thus there is some energy, $\frac{1}{4} L_{SQ} I_{SQ}^2$ associated with the SQUID loop. (This energy is a small fraction of the inductive energy stored in the tank circuit, $\frac{1}{4} L_{tank} I_{tank}^2$.)

Let us now increase the current I_{RF} supplied to the tank. Since the voltage across the tank circuit is $\omega L_{tank} I_{tank}$, it will increase linearly as the current I_{tank} increases. Eventually, the energy built up in the SQUID loop will be large enough to drive the SQUID around one hysteresis loop. That energy will be dissipated and there will be a corresponding (though slight) drop in the tank circuit current and voltage. It

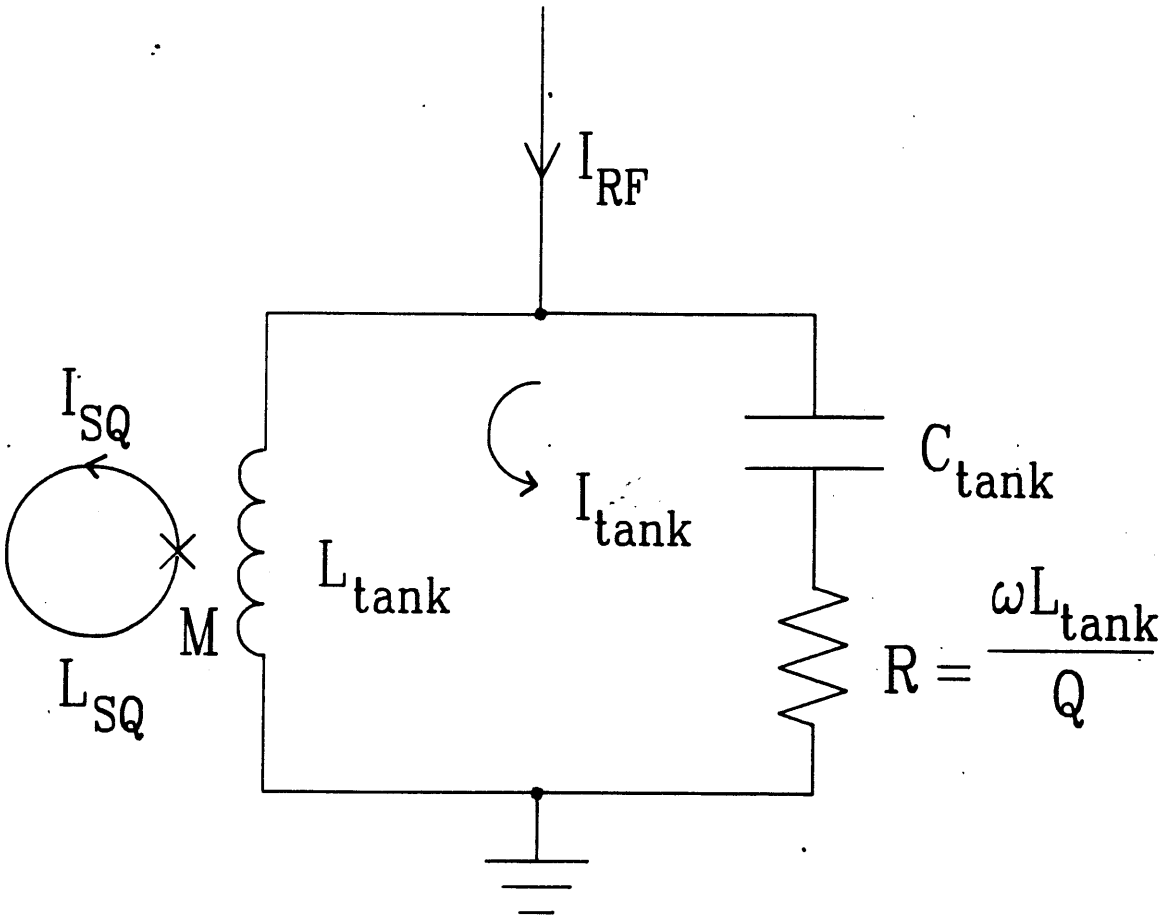

Figure III.A.1.2. Schematic diagram of typical RF excitation circuit used to pump the SQUID. As an example, our sensor uses a 19 MHz resonant circuit with $Q \approx 80$. The SQUID loop inductance is $L_{SQ} \approx 1$ nH.

then will require several cycles of RF for the tank current again to build up sufficient energy in the tank to drive the SQUID through another hysteresis loop. Therefore, increasing the drive current at this point does *not* increase the peak voltage across the tank circuit (which is pegged at its value just before the SQUID transition), but rather decreases the *time* between these flux transitions. Thus the detected peak voltage reaches a plateau. (see Figure III.A.1.3, solid line) The peak voltage remains constant until the current supplied is sufficient to drive one flux transition every half cycle of RF. Above this current, the peak voltage again increases until more flux transitions can be driven per RF cycle, and another plateau is reached, etc.

This description assumed that there was no current circulating in the SQUID loop at $I_{RF} = 0$. However, for different choices of Φ_x, there will be different values of I_{SQ} at $I_{RF} = 0$. This additional current in the ring means that the flux transitions occur at different values of I_{tank}, and thus the voltage plateau will occur at a different, lower value. (see Fig III.A.1.3, dashed line) This parametric dependence of the voltage plateau on the external, DC flux makes the RF SQUID useful as a parametric amplifier. Imagine sitting at one, specific RF level. As the external flux changes, (slowly, relative to the RF pumping frequency) the peak RF voltage will change (periodically) with the applied flux. This parametric change, as the physical basis of the RF SQUID, forms the heart of our ion detector

Of course, there are many technical problems in converting this physical principle into an ion detector. At the lowest level are the challenges in making a

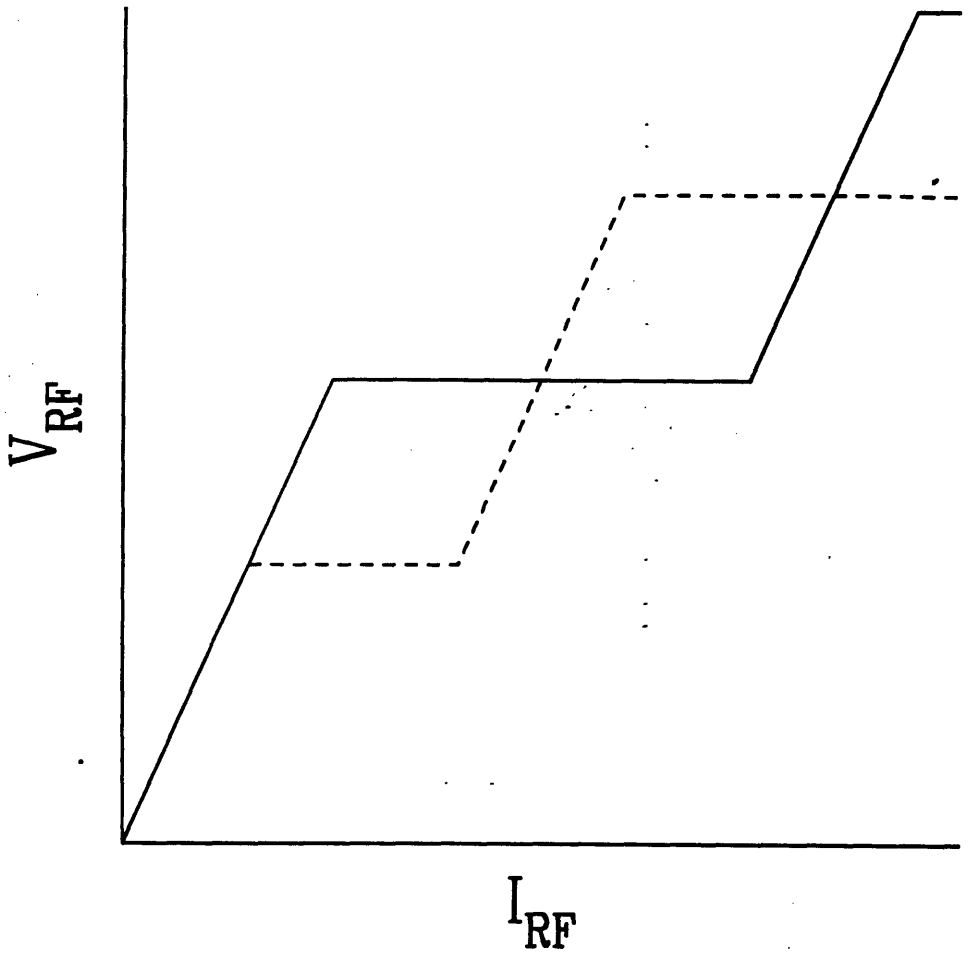

Figure III.A.1.3. Peak voltage across the tank circuit as a function of the RF current. The solid line corresponds to Φ_x such that no DC supercurrent is flowing through the loop. The dashed line corresponds to $1/4\ \Phi_0$ additional external flux.

useful flux detector: a quiet pre-amplifier, a low RF interference environment, etc. These problems are solved in our experiment with a commercial controller and SQUID sensor. At the next level, this commercial flux detector must be converted into a low-frequency current detector by adding an additional transforming linking the SQUID loop with external current sources which, in turn, induce flux changes across the SQUID. Finally, in order to operate at the ions' frequencies, we added additional feedback systems to the commercial detector at the room temperature end of the controller. I will discuss the details of transforming the commercial SQUID device into a single ion detector in the next section.

III.A.2 SQUID-Based Ion Detector

One of largest technical difficulties in this experiment was detecting the ion's motion. I gave the current induced by an ion in Equation II.A.3.4. In our apparatus, one ion, oscillating 10% of the trap size, induces just 6×10^{-15} A. Detecting this small current required understanding the noise processes of the RF SQUID detector.

There are two reasons why the SQUID noise presented theoretical difficulty. First, as Kurkijaarvi and Webb showed [KUR72, KUW73, KUR73], the presence of intrinsic fluctuations in the SQUID loop profoundly affects the V_{RF} vs. I_{RF} curve. Second, the RF SQUID acts as a parametric amplifier, and the effective input parameters (noise, impedance, etc.) of this type of amplifier are quite different from normal amplifiers. [MAR56,ROW58]

I will summarize briefly Kurkijaarvi's argument explaining the effects of noise. Fluctuations in the ring tend to spread out the value of Φ_x at which flux transitions occur. Kurkijaarvi [KUR72], viewing the transition process as a tunneling event between the stable parts of the Φ–Φ_x curve, argued the the potential barrier against tunneling decreases as Φ_x approaches the critical values, Φ_{xc}. Kurkijaarvi and Webb [KUW72] then applied this approach to a sinusoidally driven Φ_x and concluded that the V_{RF} vs.I_{RF} plateaus slope up with increasing I_{RF}. This rise occurs because of the distribution of Φ_x at the flux transition shifts with increased growth rate of Φ_x. Thus at higher I_{RF}, the increased rate of growth of Φ_x after a flux transition leads to the plateau shape.

From this insight, it took another five years for Enholm [EHN77] to incorporate the physical idea into a complete, small-signal analysis that *also* correctly treated the parametric amplifier part of the problem. His analysis was important to us for several reasons. First, he predicted that the (low frequency) input impedance of the SQUID would be negative, thus raising the possibility of oscillation. In addition, he gave predictions for various noise sources and estimated their relative sizes. We used these predictions and subsequent measurements to optimize our ion detector. We published the details of the finished detector and I include this article as the rest of this section.

This article also includes the physical details of our detector, including the construction of a high-Q, tuned circuit between the SQUID and the ion trap. This tuned circuit provides additional current gain to achieve the current sensitivity

necessary to detect single ions, but complicates the optimization of the detector. We resolved this complication in the article.

-- article --

Journal of Applied Physics, to be published.

RF SQUID Detector for Single Ion Trapping Experiments

R. M. Weisskoff, G. P. Lafyatis*,

K. R. Boyce, E. A. Cornell, R. W. Flanagan, Jr.†, D. E. Pritchard

Research Laboratory of Electronics, Massachusetts Institute of Technology

77 Massachusetts Ave, Cambridge, MA 02139

ABSTRACT

We have designed and constructed a novel superconducting detector sensitive enough to detect the axial motion of a single, trapped ion. This detector employs a tuned, superconducting transformer matched to an RF SQUID to yield an effective current sensitivity of 2.7 x $10^{-15}\ A/\sqrt{Hz}$ at 160 KHz. To optimize the detection system, we have used existing small signal equivalent circuit models, and thus have tested these models in a new regime. We include details of our superconducting circuitry and of the modifications required to stabilize the commercial SQUID controls at 160 KHz. Finally, we present typical detected signals from trapped ions.

* Current address, Physics Dept., Ohio State University, Columbus, OH 43210

†Current address, Sanders Associates, Inc, 95 Canal St., Nashua NH 03061

INTRODUCTION

Charged particle traps are used throughout the physical sciences, with applications ranging from precision measurements of elementary particles[1] to chemical analyses in commercially manufactured ion cyclotron resonance (ICR) mass spectrometers. Frequently, the signal that is detected is the small image current induced in a trap electrode as the charged particles oscillate in the trap. In many applications, then, increased detector sensitivity leads directly to improved performance. We have begun a series of experiments that involve measuring cyclotron resonance frequencies of ions in a Penning trap. Our goal is to determine the relative mass of ions to high precision; ultimately, with a single ion in the trap at a time. This paper describes an extremely sensitive detector we have developed in order to measure the small currents induced by a single, trapped ion.

We have utilized the inherent sensitivity of an RF SQUID to develop a novel detector that is capable of measuring the signal from a single ion. In our trap, this signal is a current of $< 10^{-14}$ A at 160 KHz. In developing and optimizing the detector, we used existing models of RF SQUID's in a new, higher frequency regime. We found these models account well for the SQUID's behavior.

In this paper, we also give details of our superconducting circuits, including materials and techniques. We expect that this type of detector and these techniques will be useful not only for trapped ion work but also in other applications that require

an extremely sensitive detector at sub-MHz frequencies.

THEORY

Ion Trapping

We shall briefly describe our trapped ion experiment in order to indicate the detector requirements. For further details of precision measurements with charged particles, see the recent review by Brown and Gabrielse[2].

Our goal is to measure ratios of ionic masses to very high precisions by measuring cyclotron frequencies of the ions. In order to observe the ions for a time long enough to make precision determinations of frequencies, we confine the ions in a strong magnetic field using a weak, electrostatic quadrupole trap. This configuration is known as a Penning trap. (see Fig. 1) The hyperbolic electrodes make an essentially harmonic potential in the axial direction[3]. The resulting axial motion has characteristic frequency ω_z:

$$\omega_z^2 = \frac{eV}{md^2} \tag{1}$$

As the ions move up and down in the trap, they induce oscillating image currents in the endcaps. For a parallel plate capacitor of spacing $2z_0$, we expect for an ion of charge e displaced z above the center:

$$q_{induced} = e \times \frac{z}{2\,z_0}$$

Thus, for ions oscillating at angular frequency ω, the amplitude of the induced current is:

$$i_{induced} = \frac{e\ \omega\ z}{2\ z_0} \qquad (2)$$

where z is now the amplitude of the oscillation. In a Penning trap, because of the hyperboloid geometry, $i_{induced}$ is lowered[4] by about 20%.

As a circuit element, the trapped ion may be treated as an ideal current source. Thus, in our application, the important quantity for evaluating a detector is current sensitivity. In our trap, z_0 = 0.6 cm, $\omega = 2\pi \times 160\ KHz$. For an ion oscillating with amplitude 0.1 z_0, we get 6×10^{-15} A/rms. Measuring this small current in one second thus requires a noise of less than $1.1\times10^{-14}\ A/\sqrt{Hz}$.

RF SQUIDs

An RF SQUID consists of a single Josephson junction in a superconducting ring. Its non-linear properties may be exploited to make a parametric amplifier[5]. As pointed out in a series of papers by Kurkijaarvi and Webb[6,7,8], and further incorporated into a linear, equivalent circuit by Enholm[9], the RF SQUID has some interesting noise properties. Work along these lines was also performed by Hough and co-workers[10], Gusev and Rudenko[11] and Jackel and Burmann[12].

Following Enholm, we model the input to the SQUID as a parallel inductor, L_s, and negative resistor, R_s with three noise sources. (Fig. 2) L_s is the input inductor (the current-to-flux converter), reduced a small amount by its coupling to the SQUID loop. The three noise sources are the preamplifier noise, u_p, and two correlated sources: i_{sq}

and u_{sq}, which represent the interaction of the tank current noise and the so-called intrinsic noise of the junction. Enholm provides theoretical expressions for these parameters. Typical values for the type of commercial SQUID sensor we use are found in Giffard and Hollenhorst[13,14].

We can simplify this model in two ways. In our case, the preamplifier noise swamps u_{sq}, the effective voltage noise. Thus we neglect the term u_{sq}. Also, though it would be difficult to get an exact a priori value for R_s, we can use Enholm's work and Giffard and Hollenhorst's measurements to estimate its value:

$$R_s \approx -45\ \Omega \tag{3}$$

Though this value should be taken only as an estimate, it does indicate that $\omega L_s << |R_s|$ at our signal frequencies, with L_s = 2.0 μH. Thus, except for stability questions, R_s can be neglected in what will follow.

With this small signal equivalent circuit, we can apply traditional noise theory (Robinson[15]) treating the SQUID's current noise and the equivalent noise of the preamplifier as uncorrelated current and voltage sources. We can then optimize the input circuit impedance, or, more precisely, the input circuit coupling to obtain the best current sensitivity. Note that the preamplifier noise, which we call "voltage noise" is, in fact, specified as an equivalent *current* at the SQUID's input: $i_{pre} = \frac{u_{pre}}{r}$. The sense in which it is a voltage noise is that it is independent of the source impedance, in analogy to the voltage noise of a more traditional amplifier.

Circuit

We detect the ions' motion by connecting the upper endcap to the detector. Since the total input noise of our RF SQUID is typically around $2 \cdot 10^{-11}$ A/$\sqrt{Hz}$, and the ion signal is less that 10^{-14} A. we clearly need some extra current gain. In addition, the ion looks like a high impedance current source, while the SQUID input is a low impedance inductor. To get both impedance matching and extra gain, we use a high-Q tuned transformer between trap and the SQUID. Tuned inputs have been used with SQUIDs previously[16,17], but at lower frequencies, with smaller Q's, and with less resulting current sensitivity.

We use the detector in a narrow-band fashion, driving the ions at known frequencies, and looking for their response at known frequencies. Since the ion is a current source, it must compete with the Johnson noise of the coil and the equivalent noise of the SQUID. The Johnson current noise is given by $i_R = \sqrt{\frac{4kT}{R}}$, with $R = \omega_0 L Q$, where Q is the usual quality factor of the circuit, and ω_0 is its resonant frequency. We can do no better than reducing the SQUID noise far below this Johnson noise source. This relation tells us that, ultimately, we would like to make L and Q as large as possible, and, in addition, make the SQUID's contribution to the noise negligible.

Figure 3 shows an equivalent circuit of the cryogenic part of our detector. R_s is not shown on Fig. 3 since, as mentioned earlier, we can neglect it, except for the

interesting case of oscillations. Since the preamplifier noise is given as effective current noise at the SQUID's input, we solve for the signal at that input:

$$i^2 = (i_{sm}^2 + i_R^2)\left|\frac{M}{L_2+L_s}\,\frac{1}{1-\omega^2 L_{eff}C + \frac{i\omega L_{eff}}{R}}\right|^2 + i_{sq}^2\left|\frac{L_s}{L_2+L_s}\,\frac{1-\omega^2 L_1 C + \frac{i\omega L_1}{R}}{1-\omega^2 L_{eff}C + \frac{i\omega L_{eff}}{R}}\right|^2 + i_{pre}^2$$

$$\text{where } L_{eff} = L_1 - \frac{M^2}{L_2+L_s}$$

We now calculate the optimal coupling between L_1 and L_2 to maximize the signal to noise; specifically, the optimal value for M, the mutual inductance. Weakening this coupling brings L_{eff} closer to L_1 and thus brings the minimum in the numerator of the i_{sq} term closer to canceling out the denominator. That is, weakening the coupling tends to reduce the relative contribution of the current noise of the SQUID. However, we can't weaken the coupling indefinitely, since the gain, $\frac{M}{L_2+L_s}$, decreases for $L_2 \ll L_s$. Eventually, any signal will be swamped by the preamplifier noise, i_{pre}. We will find that optimum coupling occurs when the two competing noise sources, i_{sq} and i_{pre}, have equal contributions.

We can simplify the algebra dramatically by making two approximations. First, we can neglect the imaginary part of the numerator in the i_{sq} term. The difference between L_1 and L_{eff}, when optimized, is large enough to make the real part much larger than the imaginary part near the signal resonance, $\omega_0^2 L_{eff}C = 1$. Second, we

assume that $L_2 \ll L_s$. Minimizing the total noise on resonance yields:

$$M_{opt}^2 = \frac{\gamma L_1 L_s}{Q} \qquad (5)$$

$$\text{where } \gamma = \frac{i_{pre}}{i_{sq}}$$

With this coupling, the contributions of the voltage and current noises are equal, and the total noise density becomes:

$$i_{opt}^2 = i_R^2 \frac{\gamma L_1 Q}{L_s} + 2 i_{pre}^2$$

Since $i_R = \sqrt{\frac{4kT}{\omega_0 L_1 Q}}$, we can calculate the total noise of the optimized detector, referred to the input. In our case, $i_{pre} \approx 15\ pA/\sqrt{Hz}$, $\gamma \approx 6$, and the SQUID detector noise contributes about 1/3 to the total noise power. The rest is the 4.2 K thermal noise of the tuned circuit, currently 2×10^{-15} A / $\sqrt{Hz}$, for $L_{eff} = 5.7 \times 10^{-3}$ mH and Q = 12000.

Although the current signal-to-noise ratio is the most important parameter for our experiment, it isn't the only parameter that might be interesting. For example, if we were more concerned about extracting detectable energy from the system at a fixed rate, then the noise power, $i^2\, Re(Z_{det})$ could be more relevant. In that case, it might be advantageous to detune the ions from the coil's resonance and, in addition, increase M. If the SQUID is sufficiently quiet, the noise power then can be reduced below the usual thermal limit, 4KT, without undue degradation of the signal-to-noise.

TECHNICAL DETAILS

Superconducting Transformer

We have made many inductor/capacitor pairs on the route to a high-Q combination. Since we found few details in the literature, we describe here some of the techniques we used. Rather than listing the many unsuccessful combinations of materials and techniques, we describe only what has worked best and what we believe limits us at this stage.

The primary of our present transformer has approximately 1200 turns of 4 mil, non-annealed Nb wire, insulated to 5 mils with Formvar, wound on a hollow bakelite tube, 2 cm in diameter. When placed in a Nb superconducting shield, the primary has an inductance of 5.7 mH. The secondary is just 3 turns, wound at the the end of the primary. As discussed above, its inductance is not particularly important, as long as it is much less than the the SQUID input inductance of about 2 μH. The optimized mutual inductance is 2.2 μH, but we have run with it as high at 25 μH. (see below) The capacitor is made of two cylindrical blocks of Nb, 1.7 cm in diameter, 0.3 cm high, spaced by a sheet of 1 mil Teflon, held together with Teflon threaded fasteners. Normally, the resonant frequency is stable to about 1 ppm until the liquid Helium level begins to drop below the top of the superconducting shield. The capacitance is about 100 pF.

The primary is spot-welded to Nb tabs which, in turn, are spot-welded onto the capacitor blocks. The secondary is spot welded onto the input of the SQUID sensor[18]. The transformer and capacitor are held in a machined Nb shield which is physically attached below a similar shield which houses the SQUID. This shield minimizes losses due to eddy currents. Omission of the shield or the spot-welding limits the Q to ≈2000 at 160 KHz.

Since the ion trap is in an 85000 gauss magnetic field, we need to put the detector about 100 cm above field center, where the ambient magnetic field will permit the SQUID to operate. To connect the detector to the trap, we use Teflon-insulated Cu wire, made into twisted pair and soldered onto Cu-NbTi composite wire used in superconducting magnet work. The Cu sheath of this wire is removed on one end and we spot-weld the superconducting core onto tabs located under the superconducting shield that houses the transformer. Finally the primary is spot-welded to these tabs.

In addition, since the ion trap is in a vacuum, we made a high-vacuum feedthrough of formvar insulated, Cu twisted pair, passing through SS needle tubing, sealed with Stycast. The vacuum side of the feedthrough is soldered to the trap, while the other side is soldered to the twisted pair from the detector.

With this arrangement, we routinely obtain Q's of 12000 at 160 KHz. (Without the feedthrough and twisted pair, and using a smaller, 1 mH coil, we have obtained Q's $> 2 \times 10^5$.) Most of the loss is in the twisted pair and the feedthrough that goes to

the ion trap: when tested without this line, we measure Q's of about 20000. We believe the formvar insulation is the dominant source of this loss.

Room Temperature Controls

The control electronics for the SQUID system include two separate negative feedback loops, operating in separated frequency regimes. At low frequencies, we use a commercially available SQUID control (SHE Model 30), which provides a 19 MHz pump signal and detection electronics that amplify signals at the SQUID's input. In addition, this control supplies a low frequency (< 20 KHz) "flux-locked" servo loop, which enables operation in electrically hostile environments. Since the leads running between the trap and the coil are in the strong gradients of the superconducting magnet, mechanical motion can induce significant low frequency noise.

The Model 30 controller also provides a higher frequency output, for signals beyond the locked loop. The dynamics of a hysteretic SQUID limits signals from this output (called the "HF out") to less than 1/4 of a flux quantum, ϕ_0. In addition, the gain from this output typically fluctuates ≈ 15% in eight hours. Therefore, both to stabilize the gain and increase the dynamic range, we implemented a second, higher frequency (160 KHz) feedback loop.

We take the error signal from the model 30 ("HF out"), and pass it through a tunable bandpass filter ($Q \approx 70$), followed by a variable gain, variable phase amplifier.

We return the output of this amplifier through an 800 Hz low-pass filter into the feedback path of the SQUID, effectively putting the signal at the input of the SQUID sensor. (We use the "Test In" on the Model 30.) Although our 160 KHz signal is much attenuated (since we operate far above the 3 dB point of the filter), we verified that there are no other phase shifts across the range of frequencies at which we operate.

When the feedback path is closed and the gains of the amplifiers are sufficiently high, the gain of the whole system is set by the gain of the feedback path. Since this path consists of simple circuit elements, the gain is quite stable. When our feedback loop is closed, we adjust the phase of our amplifier to create negative feedback. (When adjusted incorrectly, the system oscillates) The change in the output of the amplifier after closing the loop lets us measure the total gain of the system. Typically, we operate with an open loop SQUID gain of $\approx 3\ V/\phi_0$, and an amplifier gain of 10. The feedback gain is frequency dependent (but time independent!), but roughly $\frac{1}{4}\ \phi_0/V$ in our frequency range. With the feedback, we find much improved stability, and we can track signals up to 3 ϕ_0/pp before this external feedback loop loses lock.

We calibrated the detector in stages. We measured the current sensitivity directly, by injecting a known current into the SQUID's input without our transformer attached. We next determined the gain of the feedback path (through the "test-in") by applying a known voltage at the feedback input. We then closed our external feedback loop and measured the gain of the closed loop. This gain was consistent with the

value predicted by simple feedback theory (using the gains above), and reproducible to about 10%.

RESULTS

We have used the SQUID and transformer with several values of coupling. In preliminary runs the coupling was too tight and the detector oscillated because of the negative resistance at the SQUID. These runs also allowed us to determine the optimum choice for M, the mutual inductance between primary and secondary of our transformer. We then were able to measure the performance with optimum coupling.

Tight Coupling

When the SQUID is too closely coupled to the high-impedance LC circuit (and therefore sees an impedance which is too high), the noise on resonance will be dominated by the SQUID's effective current noise. In our initial work with the SQUID, we used M=25μH and L_2≈1.5μH. This is indeed too strong a coupling. The noise spectrum (Fig 5) shows a peak due to the current noise of the SQUID. (It is about three times larger than the theoretical 4.2 K noise of the transformer.) We can evaluate i_{sq} from this peak. Using equation (4), this peak yields i_{sq} = 2.6 pA/$\sqrt{Hz}$. With the SQUID's input shorted, on the other hand, we found $i_{pre} \simeq 15\ pA/\sqrt{Hz}$, and thus $\gamma = 6$. Using equation (6), for this primary, $M_{opt} \simeq 2.5\ \mu\ H$

In addition, with the tighter coupling, we found that the negative, real part of the SQUID's input impedance caused the detector to oscillate: an interesting demonstration of the validity of Enholm's model. We shall derive an expression for the onset of

oscillation as a function of the electrical parameters. We can transform the resistor R_s in Fig. 2, across our transformer, yielding $R_{s,eff}$, its effective resistance at the input of the detector. This resistor is in parallel with the dissipation resistance of the tuned circuit, R in Fig. 3. The system can oscillate if the parallel sum of these resistors is negative. Thus, for two resistors of opposite sign in parallel, the sign of their parallel sum carries the sign of the one with smaller absolute value. In the limit of a large primary and $-R_{sq} \gg \omega L_{sq}$ as discussed above, we find:

$$R_{s,eff} = R_s \left[\frac{L_{eff}}{M}\left(1 + \frac{L_2}{L_s}\right)\right]^2 \qquad (6)$$

where $L_{eff} = L_1 - \frac{M^2}{L_2+L_s}$, as usual.

Since the system oscillated (without our additional feedback), we conclude that $R_{s,eff} > -R$, and thus we can use equation (6) to put a lower limit on R_s. Using measured values of Q and L_{eff} to compute R, we conclude $R_s > -280\ \Omega$. When we added our external feedback loop, however, we quenched these oscillations. It was never clear whether the negative feedback itself or the noise introduced in the feedback process actually stopped the oscillations.

Finally, the pronounced asymmetry apparent in the noise peak (Fig 5) can be understood qualitatively . At the natural resonance of the primary ($\omega^2 L_1 C = 1$), with no damping, the impedance diverges, and thus no current flows in the primary. As a result, the secondary is effectively decoupled from the primary at this frequency, and

the SQUID sees only L_2, a relatively small impedance, and thus produces little current at the output due to current noise. The low frequency side of noise peak, (which occurs at $\omega^2 L_{eff} C = 1$) will be *smaller* since it is closer to the (lower) natural resonance frequency of the primary.

Optimal Coupling

When we reduced the coupling to the optimum calculated above, we saw immediate improvements: the detector stopped oscillating, the noise peak became symmetrical (Fig. 6), and the net current sensitivity was increased. (Fig. 7)

The total noise peak is 32(3) pA/$\sqrt{Hz}$ at the SQUID, which gives an effective input noise current of 2.7(3) fA/$\sqrt{Hz}$. This value is less than half the noise present with tight coupling. When we subtract the effects of i_{sq} and i_{det}, the noise remaining, referred back to the ion, is 1.8(2) fA/$\sqrt{Hz}$. This value is precisely the calculated 4.2 K noise of the resonant circuit. These errors primarily come our absolute calibration of the SQUID, with smaller contributions due to uncertainty in component values and stability in the room temperature amplifiers.

In addition, the absence of oscillation puts another limit on the negative impedance. Using equation (6) again, since $R_{s,eff} < -R$, we obtain: $R_{sq} < -10\ \Omega$.

Figure 7 also reveals a detrimental side effect of the weakened coupling: the increased current sensitivity is achieved only at the cost of decreased bandwidth. In

applications that demand wider bandwidth, a compromise toward tighter coupling may be required.

Figure 8 shows trapped N_2^+ ions. Optimum coupling was used here since the ions are driven so that they respond at the peak of the detector. To avoid capacitive feedthrough, we drive the ions at two fixed frequencies whose difference is set to the detector's peak. The ions, when present in the trap, act as a mixer. The trapping voltage then is swept, thus sweeping the ions' natural resonance frequency (see eqn. 1). In figure 8, we have driven about 10 N_2^+ ions about 1.5×10^{-2} cm, rms, and detected them with a time constant τ = 1 s. The rms noise of the detector corresponds to one ion moving up and down about 0.03 cm. Thus the detector, when optimally coupled, has sufficient sensitivity to observe the axial motion of a single ion.

SUMMARY

We have used an RF SQUID as a very sensitive ion detector suitable for single ion measurements. In order to optimize the detector, we used Enholm's equivalent circuit model which reduces the subtle SQUID noise theory into a straight forward circuit optimization problem. We have made the SQUID noise term much smaller than the 4.2 K Johnson noise of the high-Q, tuned circuit, and measured the overall current sensitivity of the detector.

We have developed techniques for constructing high-Q circuits around 160 KHz, and made simple, high frequency additions to the feedback of our commercial SQUID controls. With non-optimal coupling, we measured the current and voltage noise, which allowed us to estimate the optimal coupling. When this couping was used, we detected the noise on resonance, and found it approximately equal to the thermal noise of the tuned circuit, as predicted.

We can make improvements on this detector. By winding toroidal inductors and making lower loss feedthroughs, we can improve the Q to 50000, thus decreasing the Johnson noise by a factor of 2. If we can replace the formvar insulation on the primary, we may increase the Q even more. By making a one piece holder for the capacitor and inductor, we can improve the stability by about a factor of five. In addition, we recently obtained a more advanced SQUID control which operates at 200 MHz and which will virtually eliminate the preamplifier noise. Of course, with the

new SQUID, the optimization becomes more subtle, because the voltage and current noises are almost completely correlated. This new control unit may allow us to return to a tighter coupling, and thus increase the bandwidth of the detector. Finally, though all of our work used commercially available SQUIDs operating at 4.2 K, a similar detector could be made using high temperature superconducting technology.

We'd like to thank Phil Gould for technical assistance in the early part of the experiment. This work was supported by the NSF (PHY86-05893) and JSEP (DAAG-29-83-K-0003) and used equipment purchased with a grant from the NSF (PHY84-00992). One of up (GPL) would like to acknowledge the Donors of the Petroleum Research Fund, administered by the ACS for partial support of this work.

REFERENCES

1. R. S. Van Dyck, P. B. Schwinberg, H. G. Dehmelt, Phys Rev Lett, **59**, 26 (1987). R. S. Van Dyck, F. L. Moore, D. L. Farnham, P. B. Schwinberg, Int J Mass Spec and Ion Proc, **66**, 327 (1985), for example.

2. L. Brown and G. Gabrielse, Rev Mod Phys, **58**, 233 (1986).

3. H. G. Dehmelt, Advances in Atomic and Molecular Physics, **3**, 53 (1967).

4. G. Gabrielse, Phys Rev A, **29**, 462 (1984).

5. A. H. Silver and J. E. Zimmerman, Phys Rev, **157**, 317 (1967).

6. J. Kurkijarvi, Phys Rev B, **6**, 832, (1972).

7. J. Kurkijarvi, J Appl Phys, **44**, 3729, (1973).

8. J. Kurkijarvi and W. W. Webb, Proc Appl Supercon Conf, Annapolis, IEEE Pub. No. 72CHO 682-5-TABSC (IEEE, New York, 1972), p. 581.

9. G. J. Enholm, J Low Temp Phys, **29**, 1 (1977).

10. J. Hough, J.R. Pugh, W. A. Edelstein, W. Martin, J Phys E, **10**, 993 (1977).

11. A. V. Gusev and V. N. Rudenko, Sov Phys JETP, **45**, 637 (1977).

12. L. D. Jackel and R. A. Buhrmann, J Low Temp Phys, **19**, 201 (1975).

13. R. P. Giffard, J. N. Hollenhorst, Appl Phys Lett, **32**, 767 (1978).

14. J. N. Hollenhorst, R. P Giffard, J Appl Phys, **51**, 1719 (1980).

15. F. N. H. Robinson, **Noise and Fluctuations**, (Clarendon Press, Oxford, 1974).

16. W. B. Simmonds, W. A. Fertig, R.P. Giffard, IEEE Trans on Mag, **MAG-15**, 478 (1979).

17. S. Barbanera, et. al., J Appl Phys, **49**, 905 (1978).

18. SHE Model TSQX (low noise) SQUID sensor.

FIGURE CAPTIONS

Figure 1. The basic geometry of a Penning trap. In our trap, $\rho_0 = 0.696\ cm$, $z_0 = 0.600\ cm$.

Figure 2. Enholm's simplified model for an RF SQUID. From reference 9.

Figure 3. Equivalent circuit of the cryogenic part of our detector. The left side is the tuned transformer; the right side is the SQUID model.

Figure 4. High frequency feedback loop described in text. It is all at room temperature.

Figure 5. Detector noise at SQUID input, tight coupling, M = 25 μH. The center frequency is 157 KHz.

Figure 6. Detector noise at SQUID input, optimal coupling, M = 2.0 μH. The center frequency is 147 KHz.

Figure 7. Effective current sensitivity, for the conditions shown in Figures 5 and 6. The dashed line is an extrapolation using the model in the text.

Figure 8. N_2^+ ions. This signal comes from about 10 ions, driven 1.5×10^{-2} cm at the peak in this curve.

$$\phi(\rho,z) = V \frac{z^2 - \frac{1}{2}\rho^2}{z_0^2 + \frac{1}{2}\rho_0^2} + \text{const}$$

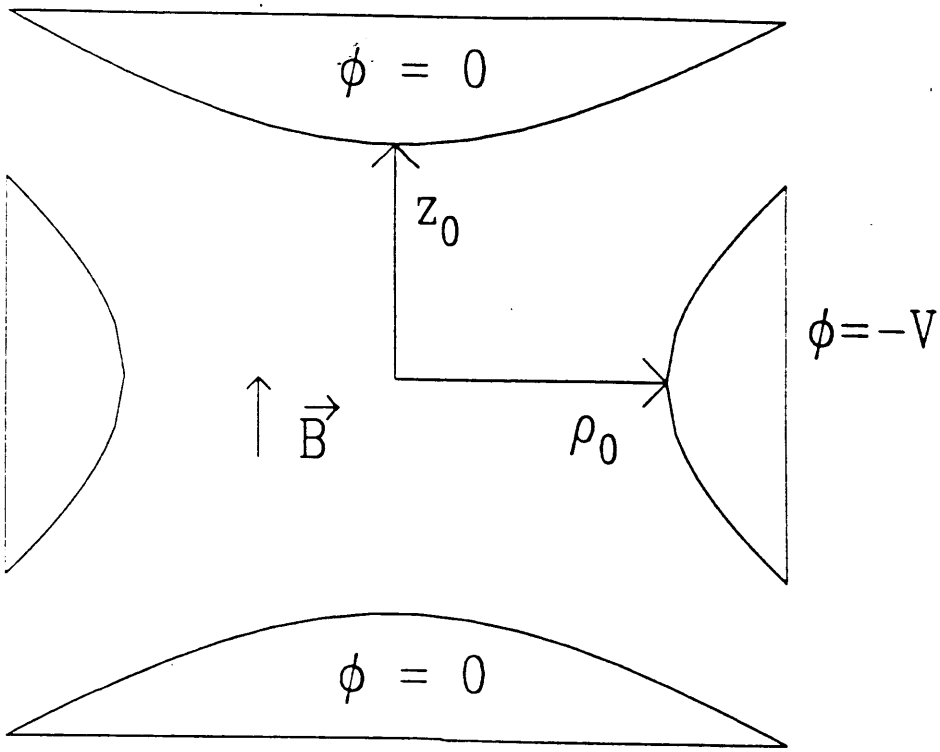

$$d^2 = \frac{1}{2}z_0^2 + \frac{1}{4}\rho_0^2$$

Fig 1

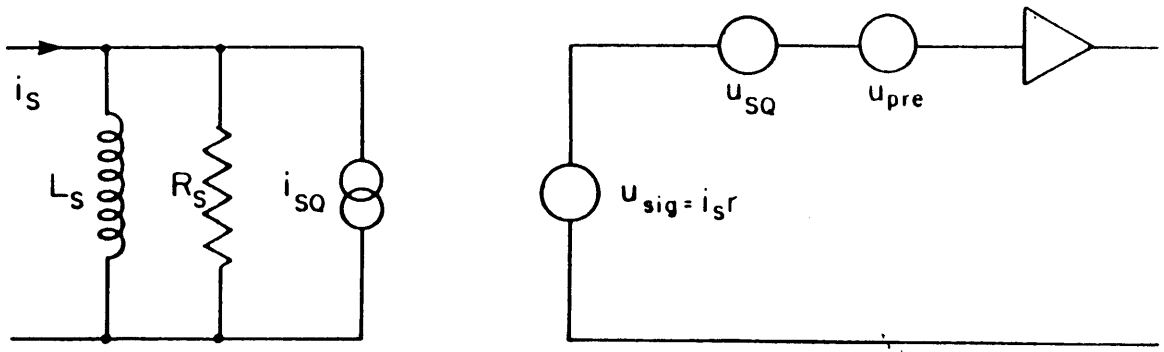

Fig 2

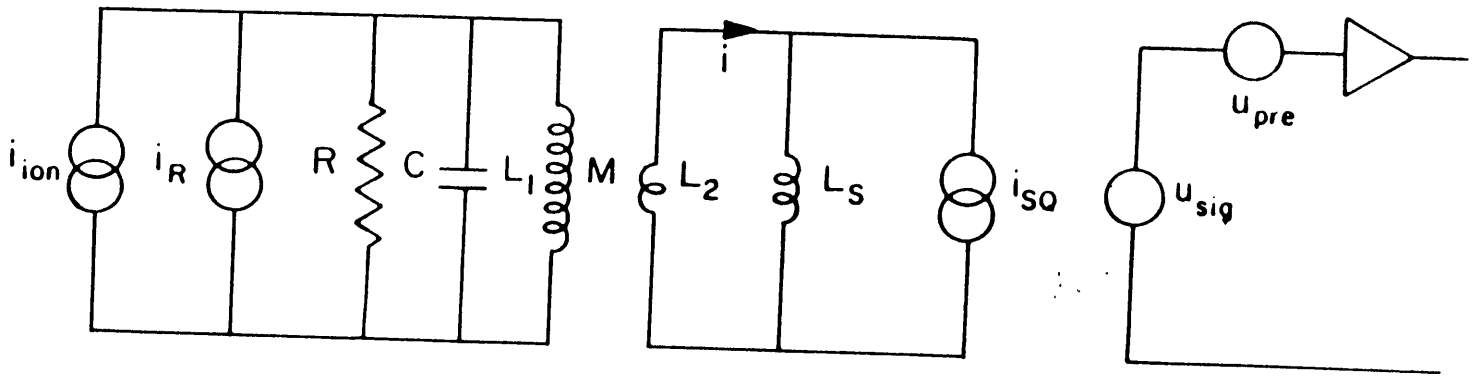

FIG 3

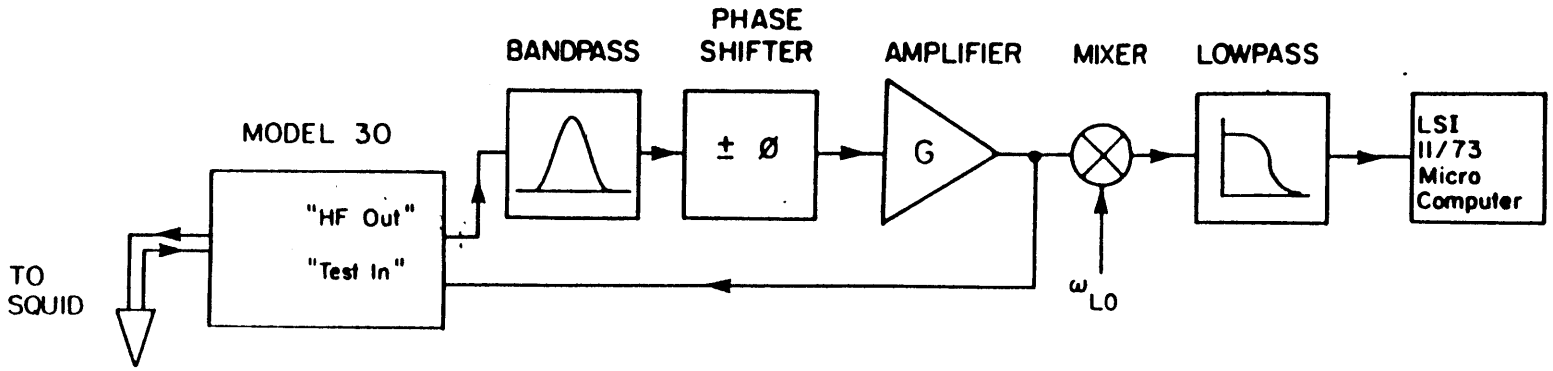
BANDPASS
PHASE SHIFTER
AMPLIFIER
MIXER
LOWPASS
MODEL 30
"HF Out"
"Test In"
± Ø
G
LSI 11/73 Micro Computer
TO SQUID
ω_{LO}
110

Fig 4

Figure 5. Tight Coupling

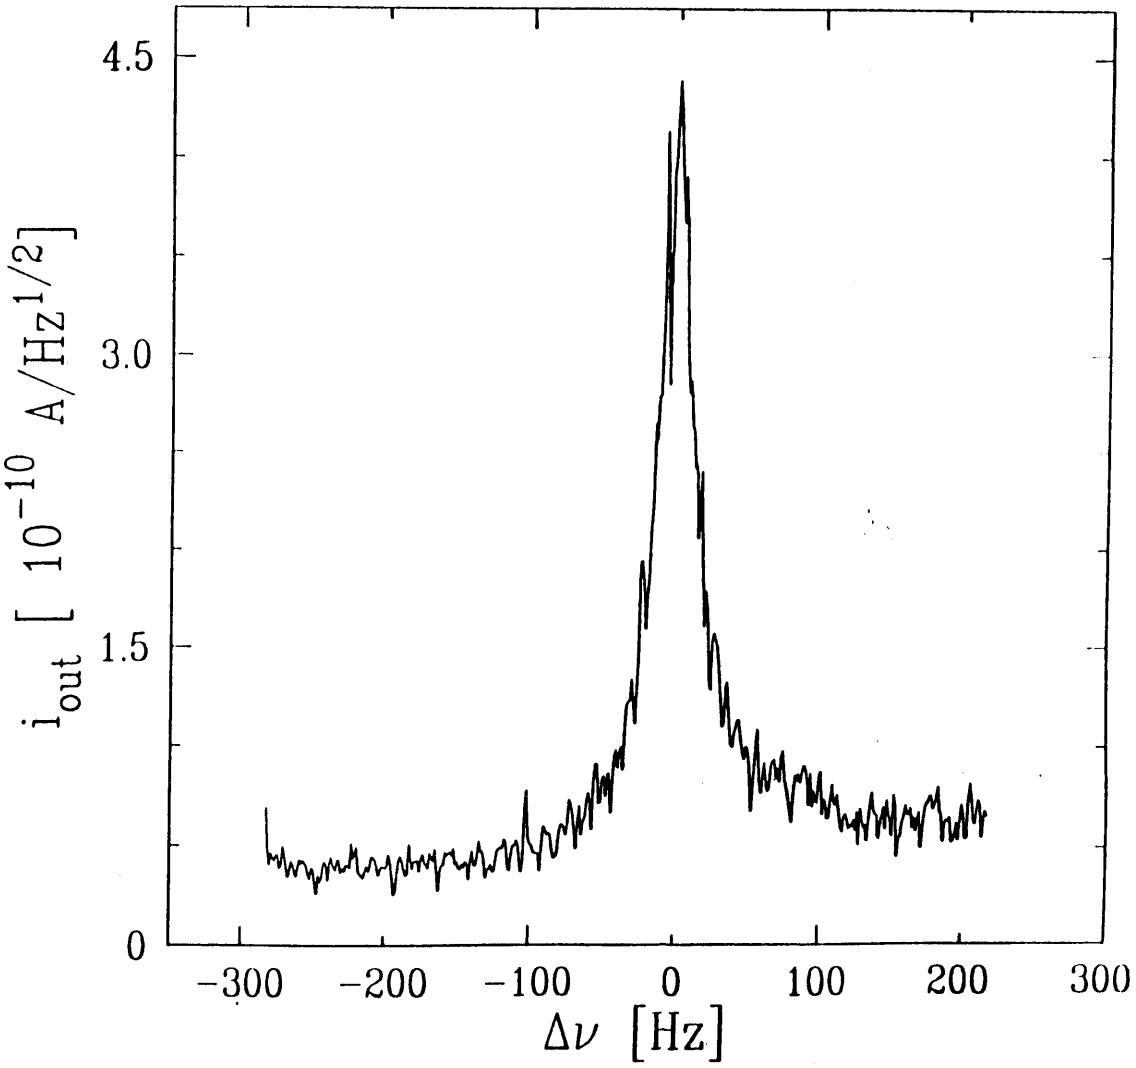

Figure 6. Optimal Coupling

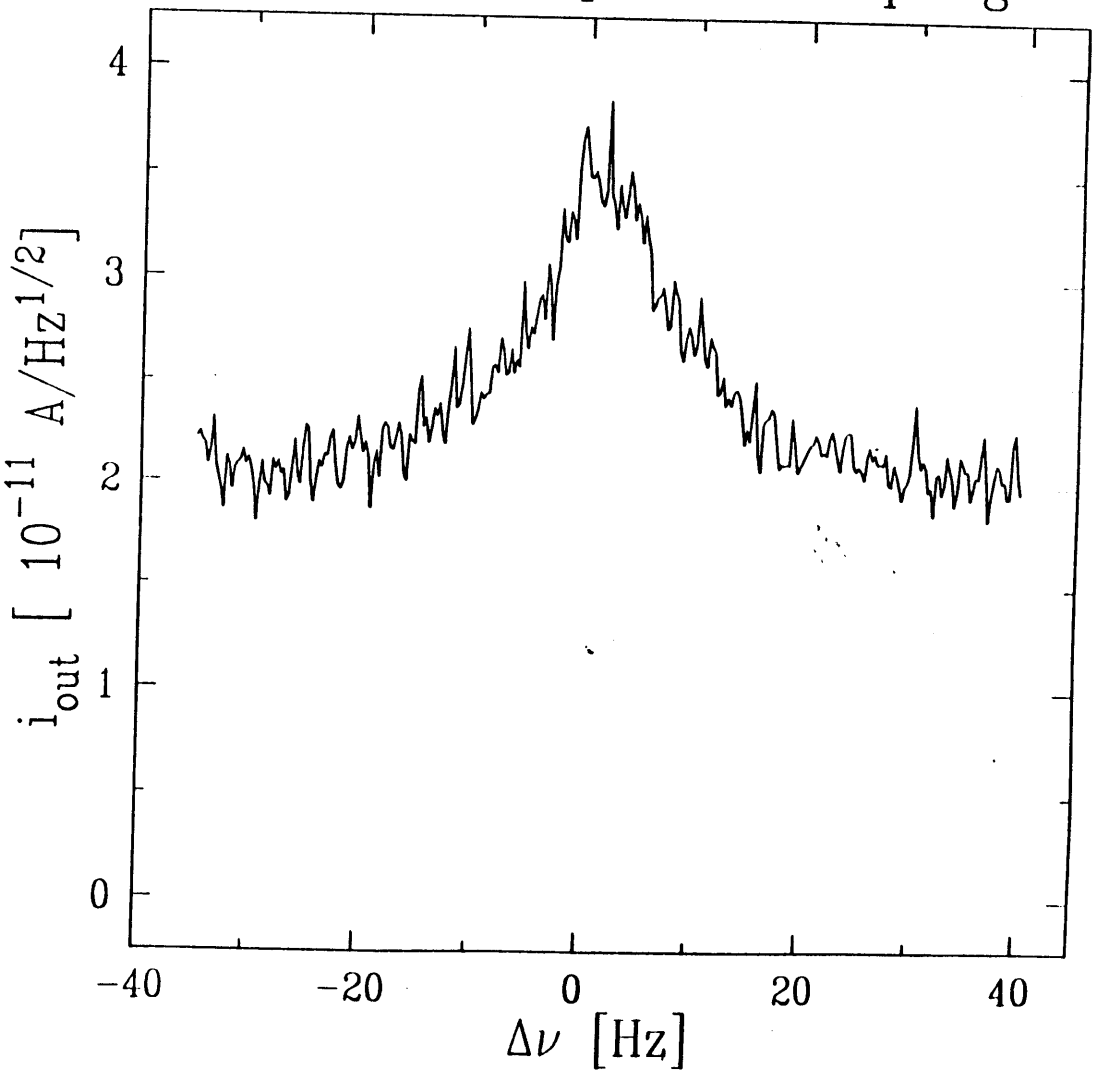

Figure 7. Noise at Ions

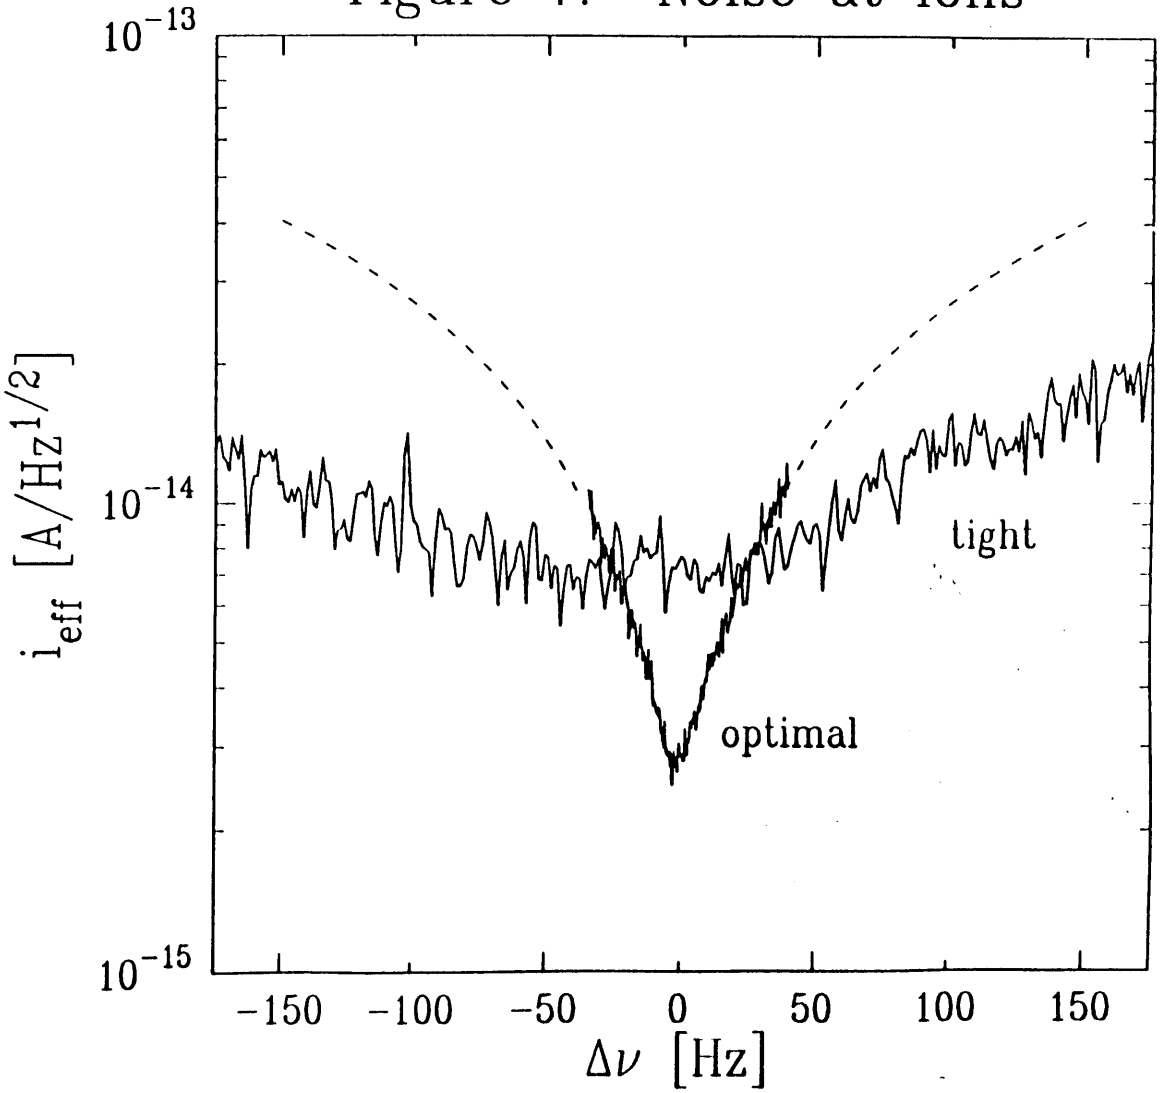

Figure 8: ~10 N_2^+ ions

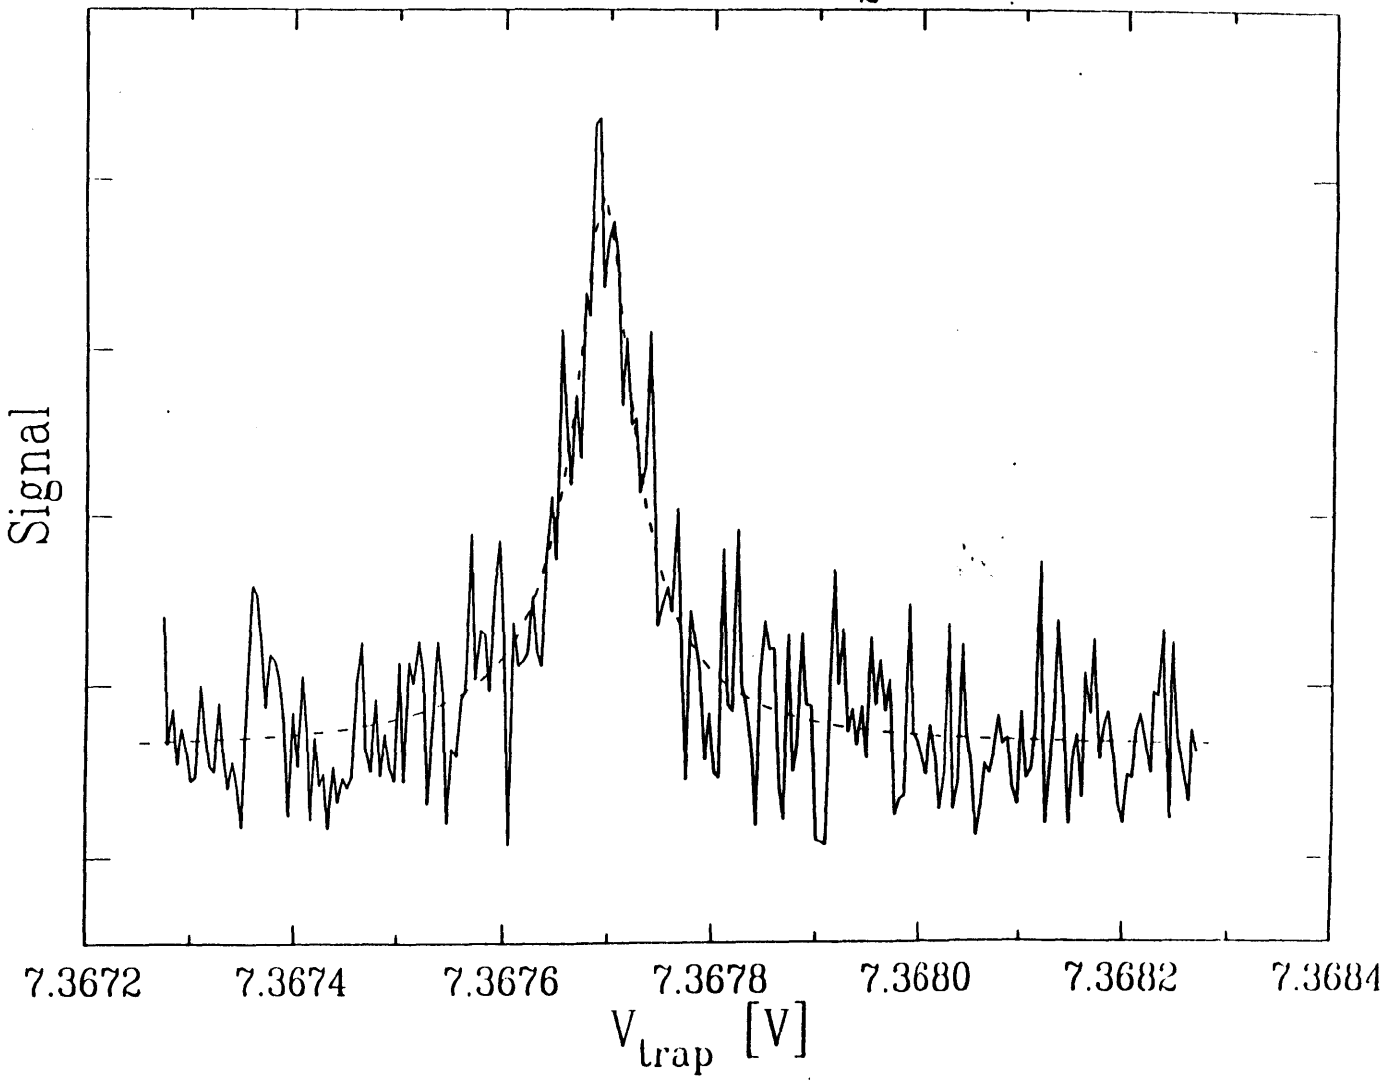

Appendix to Article

Since our submission of the Journal of Applied Physics article, we have made several improvements to the detector and understand a little better some of its limitations. We now use a torroidal inductor, wound on a teflon form, rather than the solenoidal one described in the text. This coil still has roughly 5 mH of inductance, but has a $Q > 5\times10^4$ when tested by itself. The coil and superconducting capacitor are held on a one-piece teflon form which allows less movement within the superconducting shield. We found this arrangement improved the short-term stability (~1 minute) of the resonant circuit to around 0.1 Hz.

When we wrote the article we believed that the high-vacuum feedthrough limited the Q of the detector. However, by making an identical feedthrough and testing it with the detector, we learned that the feedthrough is not the source of the dissipation. In fact, we now believe that the Macor in the trap (the machinable, ceramic spacers) provide most of the damping. We have several pieces of evidence to support this belief. First, even with the SQUID and mocked-up feedthrough attached, the Q of the superconducting resonant circuit was no less that 4×10^4. However, when the detector was hooked up to the trap, the Q dropped to about 2×10^4. In addition, each time we cooled the apparatus to liquid Helium temperature, the Q started even lower, $Q \approx 10^4$. After a day or so, the Q stabilizes at the higher value. This slow increase in Q has repeated itself every time we cycle the apparatus.

Macor, we learned, has very high AC resistivity [LAW75]. That is, even though the trap makes up a small fraction of the total capacitance of the detector, its high loss could make it responsible for the increased dissipation in the detector. In addition, this loss decreases at lower temperatures, and the *slow* decrease is consistent with the long thermalization times because the Macor spacers (which have high thermal inertia and are only press-fit onto the copper endcaps and rings) are not well-coupled to the liquid Helium thermal bath. As we shall discuss in Chapter VII, we have some plans to construct the insulators for the trap out of better materials to improve both the dissipation and the thermal equilibration problems.

With the higher Q coil, we found that we could measure the negative input resistance of the SQUID directly. (Recall that in the paper we only could set order of magnitude limits from the detector oscillation.) Because of the increased Q, the impedance of our new circuit on resonance was sufficiently large that even without our additional feedback, the detector did not oscillate. However, when we closed our external feedback loop (to stabilize the detector gain), the overall Q of the detector decreased. Using Equation 6 in the J.Appl.Phys. paper, we can determine the initial negative impedance of the SQUID from this change. Although not a particularly high-precision measurement, we get $R_{in} = -30(5)\,\Omega$, in good agreement with both theory [ENH77] and other experiments [GIH78,HOG80]. The error in this measurement comes from run-to-run instabilities in the open-loop Q; that is, the width before we stabilized it with our feedback. At its extremes, this Q varied as much a

factor of two, while the closed-loop Q was constant to better than 1% throughout the measurements.

Although we can increase the Q, it's not clear that we will want to increase it indefinitely. The increased Q lowers the Johnson noise of the tuned circuit (by increasing the impedance), but there comes a point when this increased impedance might hinder the experiment: when, as discussed in Section V.A, the detector-induced width of a single ion becomes the same order as the width of the detector. Above this impedance, the ions and detector behave as a tightly-coupled system, and shifts in the detector's resonance can cause shifts in the (observed) ion frequency. Clearly, if the detector has an impedance large enough to broaden one ion to that critical width, we will have to be very careful in the experiment. Using (V.A.1.8) and (V.A.1.5), evaluating $\mathrm{Re}Z$ on the resonance of our detector, we see that, for a single ion, we should avoid having $Q > \frac{1}{2}\sqrt{\frac{l}{L}}$. Although this is not a particularly stringent requirement for the more massive ions, for lighter particles, it can be quite severe. For example, for the Helium/Tritium comparison, we probably should avoid $Q > 5\times10^4$. Protons would present a detection problem even with the present coil.

There were also some technical points we glossed over in the article. For example, we still see occasions (albeit rare) when the detector is much noiser than usual. In addition, after several months of continuous use, we find that the voltage noise creeps up, due, we think, to relaxation of the RF shielding.[1] Also, when the

SQUID otherwise appears to be working very well (when the background voltage noise is especially low), the current noise seems elevated. We have no understanding of this relationship between the current noise and environmental conditions.

Finally, although the detector does not oscillate, because of the time-delays and gains involved in our room-temperature feedback, the overall network is underdamped. That is, like a car with bad shocks, the detector is somewhat sensitive to bounces and jolts, especially those due to environmental noise like bubbling liquid nitrogen and helium. The feedback fails to damp out these excitations very quickly, and the resulting "noise," even at 160 KHz, can be comparable in amplitude to the Johnson noise of the tuned circuit. However, these occasions are fairly rare, and usually this extra noise is not a problem.

III.B FFTs and All That

We have used the Fast Fourier Transform (FFT) extensively in this experiment. With just a little caution and understanding, the transform provides a very powerful analysis and diagnostic tool. For example, the FFT plays a central role in the pulsed modes of detection. In addition, we have used it to measure the background noise of our detector and diagnose oscillation in the SQUID system.

---

1. That is, the aluminum foil sometimes loosens and even falls off

In this section, I provide a simple introduction to the use and interpretation of the FFT as used in this experiment. I'll begin with a definition and give the simple examples of a cosine wave and white noise. Then, I will introduce some standard theory in order to explain the differences between the FFT and the full fourier transform, and use this theory to explain "windowing" and why it helps solve the problems of leakage and scalloping loss. Finally, I'll touch on some of the more subtle issues in the use of the FFT for general spectral estimation, again, within the context of our experiment.

III.B.1 Naive Spectrum Estimation

The FFT is no more or less than a speedy implementation of the Discrete Fourier Transform (DFT). It is covered in almost any introductory signal processing book (e.g., [OPS75]), as well as numerical methods books [PFT86, DSP79] and books specializing in the topic (like [BRI74]). There are many versions of the algorithm, distinguished by such labels as "decimation-in-time," "decimation-in-frequency," "Radix-2-4-8," "Complex General-N Winograd," etc., but in the end, each is just an efficient algorithm to implement:

$$X[k] = \sum_{n=0}^{N-1} x[n] e^{\frac{-2\pi i n k}{N}} \qquad \text{(III.B.1.1)}$$

In the simplest case, we take x[ ] to be an array of equally spaced samples of some signal and x[n] to be that element sampled at time $t = n / f_{samp}$, where f_{samp} is the sampling frequency. (Let me put off, for a moment, any more mathematical

treatment of what "sampling" means.) We *hope* that the $X[k]$ then somehow approximates the true fourier transform:

$$X[\omega] = \int dt\ x(t)e^{-i\omega t} \qquad \text{(III.B.1.2)}$$

where, by analogy, $\omega = 2\pi f_{samp} \times \frac{k}{N}$. I will discuss shortly just how approximately our hopes will be realized.

One quick comment on the sign convention. Physicists [PFT86] usually opt for the $e^{+i\omega t}$ in the forward transform, unlike equation III.B.1.2, partly because that convention has the correct sign for Schroedinger's equation. Electrical engineers, on the other hand, more often use the sign convention above, because this makes impedances have $+i\omega's$ in them. Since most of my detector calculations required many impedances, I've chosen to stick with the minus sign in the forward transform here, too.

Let me begin the discussion of the FFT with a familiar example: a cosine wave. In our experiment, we frequently use the FFT to detect nearly pure tones; that is, to act something like a phase detector. We can compute the FFT of a cosine wave of amplitude A:

$$x(t) = A\ cos(\omega_0 t) \qquad \text{(III.B.1.3a)}$$

$$x[n] = \frac{A}{2}e^{i\omega_0 n/f_{samp}} + \frac{A}{2}e^{-i\omega_0 n/f_{samp}} \qquad \text{(III.B.1.3b)}$$

Thus, using (III.B.1.1), we have:

$$X[k] = \frac{A}{2}\sum_{n=0}^{N-1} e^{2\pi i n\left[\frac{\omega_0/2\pi}{f_{samp}} - \frac{k}{N}\right]} + \frac{A}{2}\sum_{n=0}^{N-1} e^{-2\pi i n\left[\frac{\omega_0/2\pi}{f_{samp}} + \frac{k}{N}\right]} \qquad \text{(III.B.1.4)}$$

Assume for the moment that

$$k_0 \equiv N\frac{\omega_0/2\pi}{f_{samp}} \qquad \text{(III.B.1.5)}$$

is an integer. Then for every value of k, each sum will be zero, since we are just summing a sine or cosine over an integer number of full cycles. The only exceptions are those values of k for which the exponent is an integer times $2\pi i$: those sums contribute N. ($\sum_{n=0}^{N-1} 1 = N$). Thus:

$$X[k] = \frac{A}{2}N\delta_{k,k_0} + \frac{A}{2}N\delta_{k,N-k_0} \qquad \text{(III.B.1.6)}$$

Therefore a cosine wave input yields two spikes. In the simplest case, when the frequency of this input is less than $f_{samp}/2$, the two spikes are the positive and negative frequency components of the cosine wave. Note that the negative frequency component (i.e., the part arising from the $e^{-i\omega_0 t}$) occurs in the upper half (k > N/2) of the FFT.

Note, too, that when the signal frequency is larger than half the sampling frequency we still get two spikes. For example, when $k_0 = N + \frac{N}{4}$, corresponding to an input frequency of $1.25f_{samp}$, we get the *same* FFT as when $k_0 = \frac{N}{4}$. In fact, *all* input frequencies of the form $qf_{samp} + f_1$ where q is an integer and $f_1 < f_{samp}/2$ are indistinguishable from input signals at frequency f_{samp}. This ambiguity is called

"aliasing". Although potentially useful, it often can be misleading. In many experiments, including ours, an "anti-aliasing" low-pass filter is used preceding the sampling. This filter cuts off sharply frequencies above $f_{samp}/2$ to prevent, or at least lessen, aliasing problems.

We also measure noise using the FFT. Noise, as a stochastic process (e.g, ROB74]), is usually measured in units like power per unit bandwidth or volts / $\sqrt{Hz}$. Assume that we have an anti-aliasing filter with bandwidth Δf before the sampling. If we feed white noise (i.e., equal power per bandwidth) into the filter, what comes out when we FFT the input? The expectation of the FFT is zero, since the expectation of the input signal is also zero—it is equally likely to be positive or negative. We can compute the expectation value of the square of the FFT:

$$<|X[k]|^2> = \sum_{n,m=0}^{N-1} <x[n]x[m]> e^{-2\pi i (n-m)\frac{k}{N}} \qquad \text{(III.B.1.7)}$$

For white noise,

$$<x[n]x[m]> = <x^2>\delta_{nm} = w_n \frac{f_{samp}}{2}\delta_{nm} \qquad \text{(III.B.1.8)}$$

where w_n is the voltage noise density. (Its units are $volts^2/Hz$. The rms noise, for example, if measured on an oscilloscope would be $\sqrt{w_n \Delta f}$.)

Thus, we can evaluate the double sum (III.B.1.7) using the Kroeneker delta to do the first sum, leaving:

$$<|X[k]|^2> = w_n \frac{f_{samp}}{2} \sum_{n=0}^{N-1} 1 \qquad \text{(III.B.1.9)}$$
$$= N w_n \frac{f_{samp}}{2}$$

Notice there is no k dependence in the result, as we expect for "white" noise. Thus it appears we can use the FFT to obtain an estimate for w_n. However, if we naively attempt to do this manipulation, some problems will arise. I will return to these problems below.

What happens when the input wave, (III.B.1.3b), doesn't occur at an integer value of k_0 and, in addition, has a different phase? We can re-write (and generalize) equation (III.B.1.4):

$$X[k] = \frac{A}{2} \sum_{n=0}^{N-1} e^{i\phi} e^{2\pi i n(k_0-k)/N} + e^{-i\phi} e^{-2\pi i n(k_0+k)/N} \qquad \text{(III.B.1.10)}$$

We can sum that series to get:

$$X[k] = \frac{A}{2} \left[e^{i\phi} \frac{1 - e^{2\pi i (k_0-k)}}{1 - e^{2\pi i (k_0-k)/N}} + e^{-i\phi} \frac{1 - e^{-2\pi i (k_0+k)}}{1 - e^{-2\pi i (k_0+k)/N}} \right] \qquad \text{(III.B.1.11)}$$

If the input signal has frequency less than $f_{samp}/2$, then the values of k for which X[k] is large will be those values that make the denominator of the first term in (III.B.1.10) small; that is, for k near k_0. (In this case, the denominator of the second term is *never* small.) Near k_0, assuming $k - k_0 \ll N$:

$$X[k] \approx \left[\frac{A}{2} N \right] e^{i\phi} \frac{1 - e^{2\pi i \varepsilon}}{2\pi i (k_0 - k)} \qquad \text{(III.B.1.12)}$$

where $\varepsilon \equiv$fractional part of k_0. The part of (III.B.1.12) in parentheses is the constant

from (III.B.1.5). The other part of (III.B.1.12) shows that the delta function has spread over other frequencies, but peaks near k_0. This spreading is called "leakage" in the literature. In addition, the peak value now is:

$$X[k]_{peak} \sim \frac{1 - e^{2\pi i \varepsilon}}{2\pi\varepsilon} \qquad \text{(III.B.1.13)}$$

Thus, for $\varepsilon = ½$, the worst case, the peak is lowered by $\frac{2}{\pi} \approx 0.64$. This lowering of peak values with small changes in input frequency is called "scalloping." [HAR78] Another problem, especially for more precise measurements, is that scalloping is phase-dependent. Although I won't demonstrate this effect explicitly, it certainly seems reasonable from equation (III.B.1.12). If we look at $|X[k]|^2$, the interference term will have a $\cos(2\phi)$ in front, and, it turns out, depend on ε.

If we know the signal's frequency ahead of time, we can make scalloping and leakage quite small by arranging k_0 to be an integer, either by manipulating the signal frequency (if we can) or by changing the sampling frequency. The FFT then becomes exactly analogous to phase detection, and any scalloping or leakage will come about because of timing errors. However, in most cases that we use the FFT, we don't know the frequency ahead of time (otherwise we *would* use a lockin!), and thus we need some way to handle these problems. The usual handle chosen is called "windowing." However, to explain why windowing ameliorates these problems, I need to introduce a bit more theory.

III.B.2 The FFT, Fourier Transform, and Sampling

In this section, I describe the FFT of a sampled process in more mathematical detail. This description is not meant to be rigorous, but merely to provide sufficient theory to understand some of the signal processing choices made in the experiment. This discussion follows the illustrations of Brigham, but the words of Oppenheim and Schafer, but similar descriptions can be found in many digital signal processing texts. It will assume an understanding of "normal" fourier transforms and their properties.

This treatment uses delta-functions to represent the discrete and finite elements of the DFT, and thus lets us make contact, through multiplication and convolution in the time and frequency domains, with the continuous signals, and thus with our intuition about the continuous transforms.

The sampling process can be described mathematically as a sum of δ-functions whose amplitude is the value of the (continuous) process evaluated at the time when the δ-function is non-zero. That is, if x(t) is the continuous process, and we sample every t_0 seconds, then:

$$x^*(t) = \sum_n \delta(t - nt_0) x(nt_0) \qquad \text{(III.B.2.1)}$$

is the sampled process. Thus sampling can be thought of as multiplication by the δ-function stream, $\Delta(t) = \sum_n \delta(t - nt_0)$. (See Fig III.B.2.1a).

Consider now the fourier transforms of $x(t)$, $\Delta(t)$, and $x^*(t)$. Let $\tilde{X}(\omega)$ be the

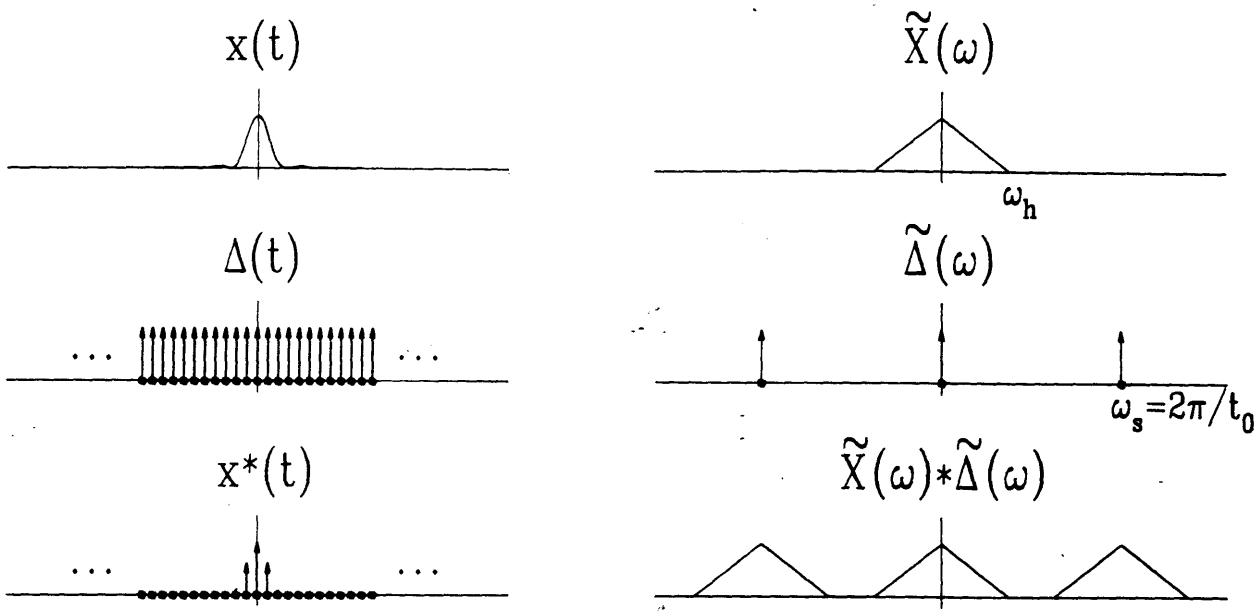

Figure III.B.2.1. The left (a) and right (b) are the time and frequency domain representations, respectively. The top graphs are for a simple input signal. The middle graphs represent the sampling process. The bottom graphs show the *sampled* results.

fourier transform of x(t). Also, let the fourier transform of $\Delta(t)$ be:

$$\tilde{\Delta}(\omega) = \int_{-\infty}^{\infty} dt\ e^{-i\omega t} \sum_n \delta(t - nt_0) = \sum_n e^{-i\omega n t_0} \qquad \text{(III.B.2.2)}$$
$$= \frac{2\pi}{t_0} \sum_n \delta(\omega - \frac{2\pi n}{t_0})$$

Since $x^*(t)$ is the product of $\Delta(t)$ and x(t), its fourier transform is given by the convolution $\tilde{X}(\omega) * \tilde{\Delta}(\omega)$. Thus the fourier transform of the sampled x(t) is the periodic replication of the fourier transform of x(t), repeated at intervals of $\frac{2\pi}{t_0}$. (see Fig III.B.2.1b).

Note that if $\tilde{X}(\omega)$ is zero outside of $\pm\frac{\pi}{t_0}$, there will be no overlapping of the successive replications, and we could reproduce the full x(t) from the section of $\tilde{X}(\omega) * \tilde{\Delta}(\omega)$ between $\pm\frac{\pi}{t_0}$. This fact is the essence of the Nyquist sampling theorem [OPS75] which states, approximately, that no information is thrown away if one samples at a rate greater than twice the highest frequency component of the input signal. However, if there are frequencies greater than $\omega_{sample}/2$ (see Fig III.B.2.2), then the input signal cannot be reconstructed uniquely. This overlapping is aliasing, discussed in the previous section for a cosine wave.

This discussion assumes that we have the full time evolution of the input signal at our disposal. On the other hand in any real experiment we can only sample for a finite time. We represent this finiteness by multiplying $x^*(t)$ by a time window, w(t).

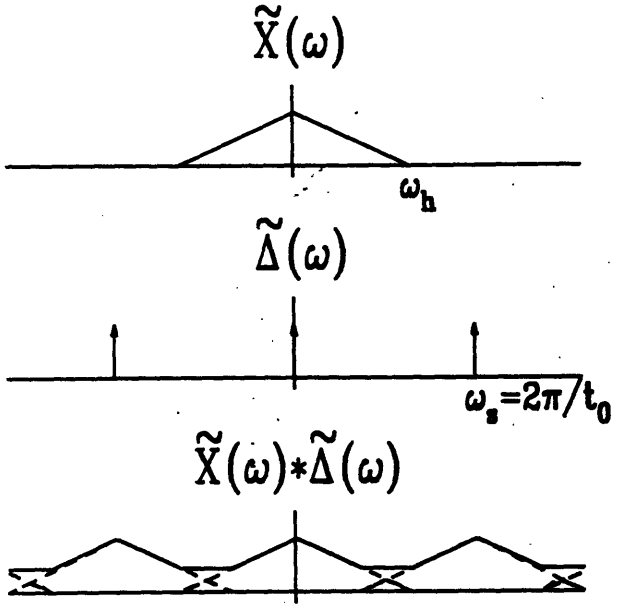

Figure III.B.2.2. Aliasing. When the signal to be sampled has frequency components greater than $\omega_s/2$, the fourier transform of the sampled signal will *not* be simple replicas of the fourier transform of the original signal. Instead, there will be significant overlap. (This is an illustration of the Nyquist sampling theorem.)

For simplicity, consider a window centered at t=0, which is non-zero for a total time T. (see Fig III.B.2.3a) Call the resulting time domain function y(t). Its transform, $\tilde{Y}(\omega)$, is the convolution of $\tilde{X}(\omega)*\tilde{\Delta}(\omega)$ with the fourier transform of the window, $\tilde{W}(\omega)$. In the case of the rectangular window, $\tilde{W}(\omega)$ is a sinc function, $\tilde{W}(\omega) \sim \frac{\sin(\omega T/2)}{\omega T/2}$.

The final step to the FFT is to sample the frequency values of the resultant transform, $\tilde{Y}(\omega)$. Just as its time analogue, the frequency sampling is expressed mathematically by multiplying $\tilde{Y}(\omega)$ by a stream of δ-functions, $\frac{1}{2\pi}\sum_k \delta(\omega - \frac{2\pi k}{T})$. For convenience, we can consider the coefficients in front of the δ-functions, $X[k=\frac{\omega T}{2\pi}]$ as an array. This array, then, I call the DFT of our original signal, x(t). Expressing x(t) as an array, and shifting the origin of time by $\frac{T}{2}$, we get my original definition of the FFT, Equation (III.B.1.1).

The sampling of $\tilde{Y}(\omega)$ in the frequency domain becomes, in the time domain, convolution with the stream $\sum_m \delta(t-mT)$; that is, the periodic replication of the windowed input data. In words, then, the DFT is the fourier transform of the sampled, windowed, periodically replicated input data. (See Fig III.B.2.3b)

Leakage and Windows

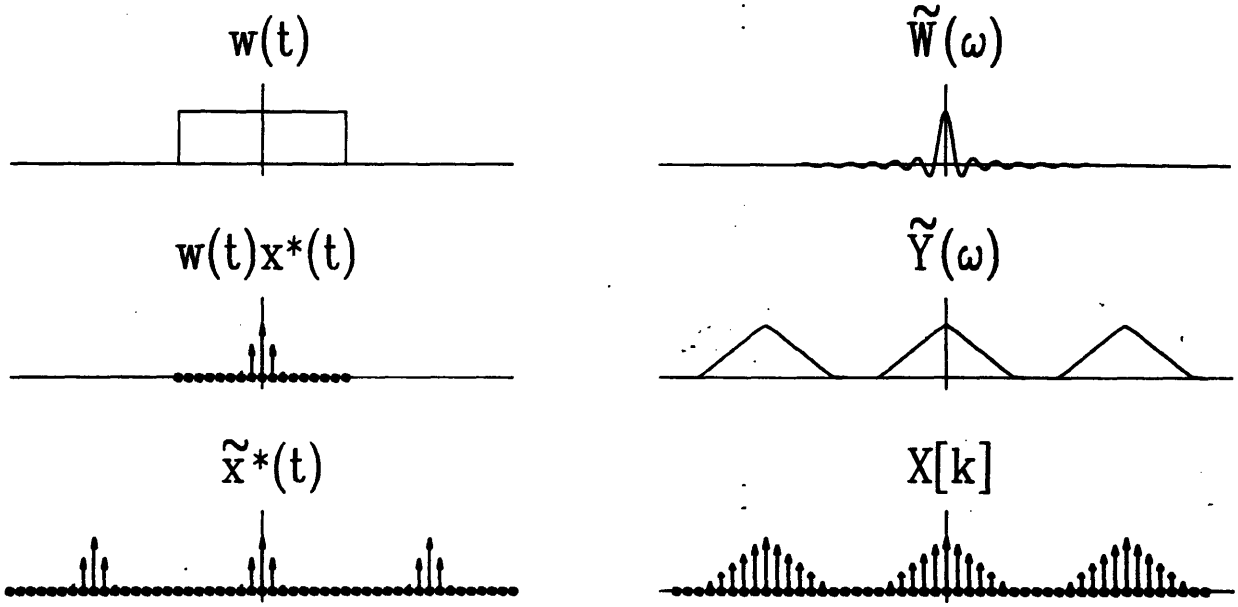

Figure III.B.2.3. The left (a) and right (b) are the time and frequency domain representations, respectively. The top graphs represent the finiteness of our sampling interval. The middle graphs incorporate the window into the results in Figure III.B.2.1. The bottom graphs show the results of sampling in the frequency domain.

This formalism lets us understand leakage and helps explain why various windows can be useful. Imagine again sampling a cosine wave, x(t) = A cos(ω_0t). (I've shifted the sampling window to start at t=0 for convenience.) Since the DFT is the fourier transform of the periodic replication of the windowed original, we see that only when x(0)=x(T) is the periodic replicate precisely the original wave. (Fig III.B.2.4a) This boundary condition is $\omega_0 T = 2\pi k$, where k is an integer. Since T = N/f_{samp}, this requirement is the same condition as (III.B.1.5): there is no leakage.

For the cosine wave, $\tilde{Y}(\omega)$ (the fourier transform of the sampled, windowed input) is the periodic replication of the sinc function. We can ignore the replications and focus our attention on the neighborhood of ω_0 when we sample $\tilde{Y}(\omega)$ to get y[n]. If $k_0 \equiv \frac{\omega_0 T}{2\pi}$ is an integer, then the samples all fall on the nodes of the sinc function, except for the one sample at k_0. If, on the other hand, k_0 is not an integer, (Fig III.B.2.4b) the samples fall on different parts of the sinc, and we get leakage.

But there is no need to use the rectangular window to do the apodization. In fact, almost any function that is zero outside the sampling interval can be used. As one (cynical?) observer noted:

> There is a lot of perhaps unnecessary lore about choice of a window function, and practically every function which rises from zero to a peak and then falls again has been named after someone. ([PFT86], p. 425)

There are, indeed, many windows, and, it seems, almost as many figures-of-merit to

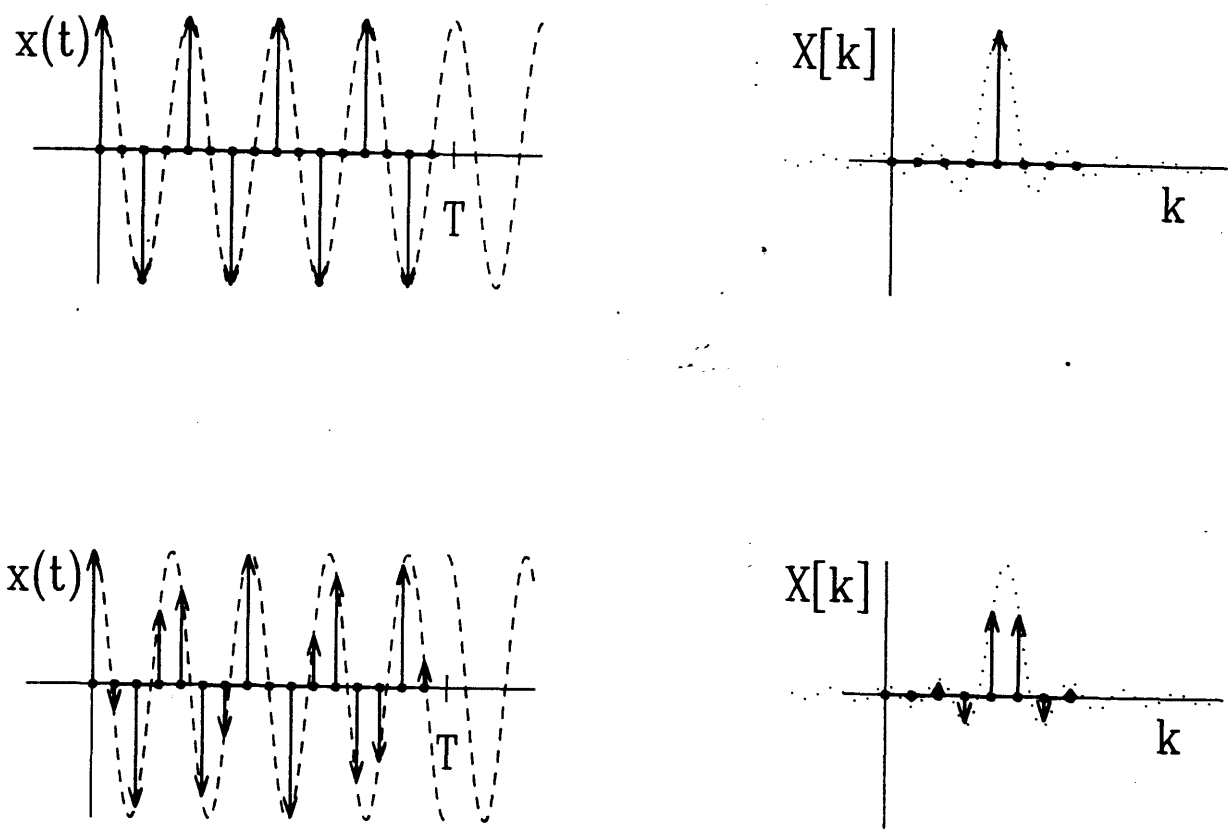

Figure III.B.2.4. Leakage. The top (a) shows a cosine wave that "fits" neatly into the sampling interval. The samples of the DFT (top, right) are non-zero at only one point. The bottom (b) shows a cosine wave that is π radians away from fitting into the sampling interval. The samples of the DFT (bottom,left) are non-zero throughout the spectrum.

judge these windows. (see, for example, [HAR78]).

It is easy to see how windows reduce leakage. If the fourier transform of the window falls off quickly in the side lobes (and almost any smooth function falls off faster than a rectangular window), then the contributions from the nearby samples of $\tilde{Y}(\omega)$ will be smaller.

For example, the "Hanning" (or "Tukey") window is:

$$w(t) = \begin{cases} \frac{1}{2} - \frac{1}{2}\cos\frac{2\pi t}{T} & 0<t<T \\ 0 & t<0,\ t>T \end{cases} \qquad \text{(III.B.2.3)}$$

and fourier transforming yields,

$$\tilde{W}(\omega) \sim \frac{\sin\frac{\omega T}{2}\left[\frac{2\pi}{T}\right]^2}{\omega(\omega-\frac{2\pi}{T})(\omega+\frac{2\pi}{T})} \qquad \text{(III.B.2.4)}$$

(see Figure III.B.2.5) Since the sides lobes now fall off like ω^3 instead of ω, leakage is dramatically reduced. However, notice that the main lobe is now *broader*. Thus we have achieved the narrowing of the broadest at the expense of broadening the most narrow: even if k_0 is an integer, the peak in the FFT will extend over three bins.

Another effect of using a Hanning window is a lessening of scalloping: when ε=½ (see III.B.1.12), the coherent peak is lowered to $\frac{8}{3\pi}\approx 0.84$, an improvement of more than a factor of two over the rectangular window.

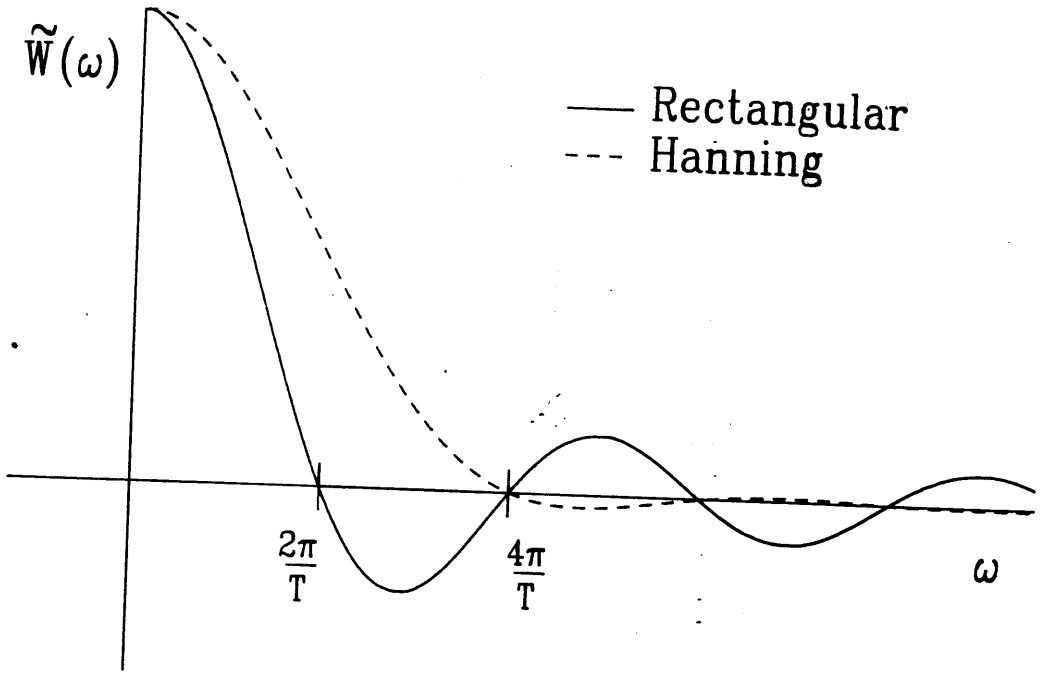

Figure III.B.2.5. Windowing to improve leakage. By using a better apodization function (the dashed line), the side-lobes of the fourier transform are much smaller, hence leakage is reduced.

In addition, a window affects noise differently than in affects a pure tone. In general, we can do calculations similar to (III.B.1.4) and (III.B.1.9) to asses the impact of the windows. [HAR78] An integer k_0 tone will have its amplitude peak reduced by a factor:

$$c_1 = \frac{1}{N}\sum_{n=0}^{N-1} w[n] \qquad \text{(III.B.2.5)}$$

relative to the same tone analyzed with a rectangular window. For a Hanning window, c_1=0.50. In addition, a general window will reduce $<|X(k)|^2>$ for white noise by a factor:

$$c_2 = \frac{1}{N}\sum_{n=0}^{N-1} w[n]^2 \qquad \text{(III.B.2.6)}$$

relative to a rectangular window. For a Hanning window, c_2=0.375.

Notice that, for any window, $c_2 > c_1^2$. Thus, for a pure tone with additive white noise, there will always be a *decrease* in the signal-to-noise of the integer-k_0 tones. For example, using a Hanning window, we expect about a factor of 1.2 decrease in the amplitude S/N. However, this decrease is more than offset by the decrease in scalloping loss for the non-integer k_0 tones, and by the ability to distinguish weak tones near stronger ones. Besides, as mentioned before, if we can always arrange k_0 to be an integer, a phase-detector is probably the more appropriate instrument.

III.B.3 More Spectrum Estimation

Finally, let me discuss the FFT as a tool for spectrum estimation one last time. The goal of this discussion is to explain the averaging scheme we use in the experiment, and *not* the general use of FFT's in modern signal's applications. I will avoid a rigorous discussion that makes use of the autocorrelation function, and instead, motivate rather than prove the commonly used techniques which we, too, have exploited. A drawback of this approach is that it leaves "power spectrum" undefined. For a proper discussion of power spectra, see OPS or PAP or ROB.

The simplest estimate of the power spectrum is the "periodogram:" [OPS75] the norm-squared of the FFT. I have already discussed the fact that the FFT is a "biased" estimate; that is, it doesn't converge to the power spectrum regardless of the number of data points we use in the FFT. Leakage is an example of this bias. It turns out that the FFT is also an inconsistent estimate: additional points do not force the variance of the estimate to zero.

We can illustrate this inconsistency by continuing the analysis of white noise which ended with Equation III.B.1.9. We already calculated $<|X(k)|^2>$ for white noise. What is the variance associated with this estimate? In general:

$$\mathrm{var}(z) = \mathrm{E}[z^2] - (\mathrm{E}[z])^2 \qquad \text{(III.B.3.1)}$$

where E[] means "estimate." Taking z to be $<|X(k)|^2>$:

$$\mathrm{E}(z^2) = \sum_{n,m,p,q} <x[n]x[m]x[p]x[q]> e^{-2\pi i(n-m)k/N} e^{-2\pi i(p-q)k/N} \qquad \text{(III.B.3.2)}$$

For gaussian processes, there is a standard factorization of the 4-point expectations

value [BRG86]:

$$\langle x[n]x[m]x[p]x[q]\rangle = \langle x[n]x[m]\rangle\langle x[p]x[q]\rangle + \text{the 2 other permutations of nmpq} \qquad \text{(III.B.3.3)}$$

and, substituting (III.B.3.3) into (III.B.3.2), using (III.B.1.8):

$$\langle |X[k]|^4\rangle = \sum_{n,m}\langle x^2\rangle^2\left\{2 + e^{-4\pi i(n-m)k/N}\right\} \qquad \text{(III.B.3.4)}$$

The second term in the braces is zero except for m=n or m=$\frac{N}{2}$+n, and thus, when summed, that term only contributes ~ N, whereas the first term contributes ~ N^2. We thus neglect the second term, and write:

$$\langle |X[k]|^4\rangle = 2N^2\langle x^2\rangle^2 \qquad \text{(III.B.3.5)}$$

We plug the two estimates III.B.3.6 and III.B.1.9 into the variance expression, III.B.3.1:

$$\begin{aligned} \text{var}|X[k]|^2 &= 2N^2\langle x^2\rangle^2 - N^2\langle x^2\rangle^2 \\ &= N^2\langle x^2\rangle^2 \\ &= \langle |X[k]|^2\rangle^2 \end{aligned} \qquad \text{(III.B.3.6)}$$

Thus, the variance is *always* the square of the mean, no matter how many points we take. Doubling the number of points, for example, appears to fail to improve the estimate of $|X[k]|^2$ at all!

The resolution of this paradox is quite simple. Although we haven't improved the estimate of any one particular X[k], were we to average over the whole spectrum, then the estimate indeed would improve by the expected factor. That is, the information obtained with the additional points provides finer frequency resolution

rather than improving the values of the existing points. Averaging, it turns out, is the generally approved method for improving variance in FFT's.

For example, to get an improved estimate of the spectrum of the detector noise in our experiment, we used the "Welch procedure" [WEL67] of averaging windowed spectra. This procedure is simple to understand. Assume a total of M points of data are recorded. Rather than doing the M-point FFT of the whole data set, the M points are broken into a series of N-point FFTs. (M=PN) The first N points are multiplying by a window (in our case, a Hanning window as mentioned above) and fourier transformed. The resulting periodogram is stashed away. The window then is shifted over by N/2 points; that is, the next N points are taken from point $\frac{N}{2}$ to point $\frac{3N}{2}$, windowed and transformed. The resulting periodogram is added to the one obtained from the last step.

This procedure is continued until the whole data set has been transformed: a total of 2P-1 N-point FFTs. The windows are overlapped to avoid "throwing away" good data. Since windows tend to weight data in the middle move heavily than data at the edges, overlapping by 50% takes the data at the edge of one iteration to the middle of the window during the next. It can be shown (see [OPS75], section 11.4.3 for example) that this procedure eliminates inconsistency: the variance of $<|X[k]|^2>$ *does* decrease like the number of periodograms averaged, P. At constant N (constant FFT size), increasing points decreases variance like the inverse of the total number of

points. In addition, the windowing has the desired effect of decreasing biases (like leakage) as discussed above.

A final word about FFTs. With the relative ease of non-linear fitting using standard software packages, it becomes reasonably simple to fit periodograms to expected shapes, making it possible to reduce or eliminate the effects of bias, split lines, determine damping, etc. from *un*windowing transforms like (III.B.1.4). In these cases, the leakage is used to help find where, to much smaller than a single frequency bin, a signal occurs. However, fitting every signal would be excessively time consuming, and we use this procedure rarely, for example, to measure the decay time and frequency of a pulsed ion.

CHAPTER IV

APPARATUS

In this chapter, I describe our ion trapping apparatus in greater detail. Much of the overall structure and environment for our experiment was described in a previous dissertation [FLA87]. Therefore, I will detail here only substantial changes since that time. Before describing these improvements, though, in Section IV.A I will briefly summarize relevant parts of the apparatus.

We made several changes so that we could measure the radial motions and shift the ions within the trap. In Section IV.B, I discuss modifications to the Penning trap. A schematic for the low-temperature, driving and filtering circuits will be presented in Section IV.C.

We also modified the DC voltage source to improve its stability, flexibility and to make computer-control easier. (Section IV.D) By replacing the Mercury cells (our old voltage source) with precision voltage references powered by rechargeable gel-cells, we also made this source more convenient and reliable.

In Section IV.E, I will describe the computer system and discuss some of our signal processing. We made several computer-based instruments which became indispensable for the experiment. We used one type, based on a software synchronous detector—essentially a computer-controlled and -implemented lock-in amplifier—in all the CW experiments. I will also describe our coherent transient-averager which became invaluable for improving the signal-to-noise in the pulsed experiments.

I will conclude this chapter with a description of how we prepared the apparatus for a run; starting with alignment, continuing though sealing and cooling, finishing with detecting and tuning-up the ions' resonance. This final section will provide a link to the results in the rest of the thesis.

IV.A Overview of the Experiment

Figures IV.A.1, IV.A.2, and IV.A.3, (reproduced from [FLA87]) present a quick sketch of the apparatus. The penning trap (modifications discussed further in Section IV.B) hangs at field-center of an Oxford 360/89 persistent-mode, superconducting, 8.455 Tesla NMR magnet. The trap was attached at the end of a two meter-long, 3/4" O.D., thin-wall stainless steel tube. Directly above the trap (in Figure IV.A.3) are several copper platforms on which sit various capacitors, inductors and resistors. These components filter the various trapping, shifting, compensating and driving potentials. This end of the apparatus was sealed within an OFHC copper can to complete the high-vacuum enclosure.

We attached the SQUID sensor and the superconducting tuned circuit about 75 cm above the top of the Copper can. Both the sensor and the coil were held within individual Niobium cans. Since these shields, the coil, and the SQUID loop all are made of pure Niobium (a Type-I superconductor with a critical field of about 2000 gauss [CRC]) we needed to keep the detector above the strongest parts of the magnetic field. (It presently resides in a region of about 200 gauss.)

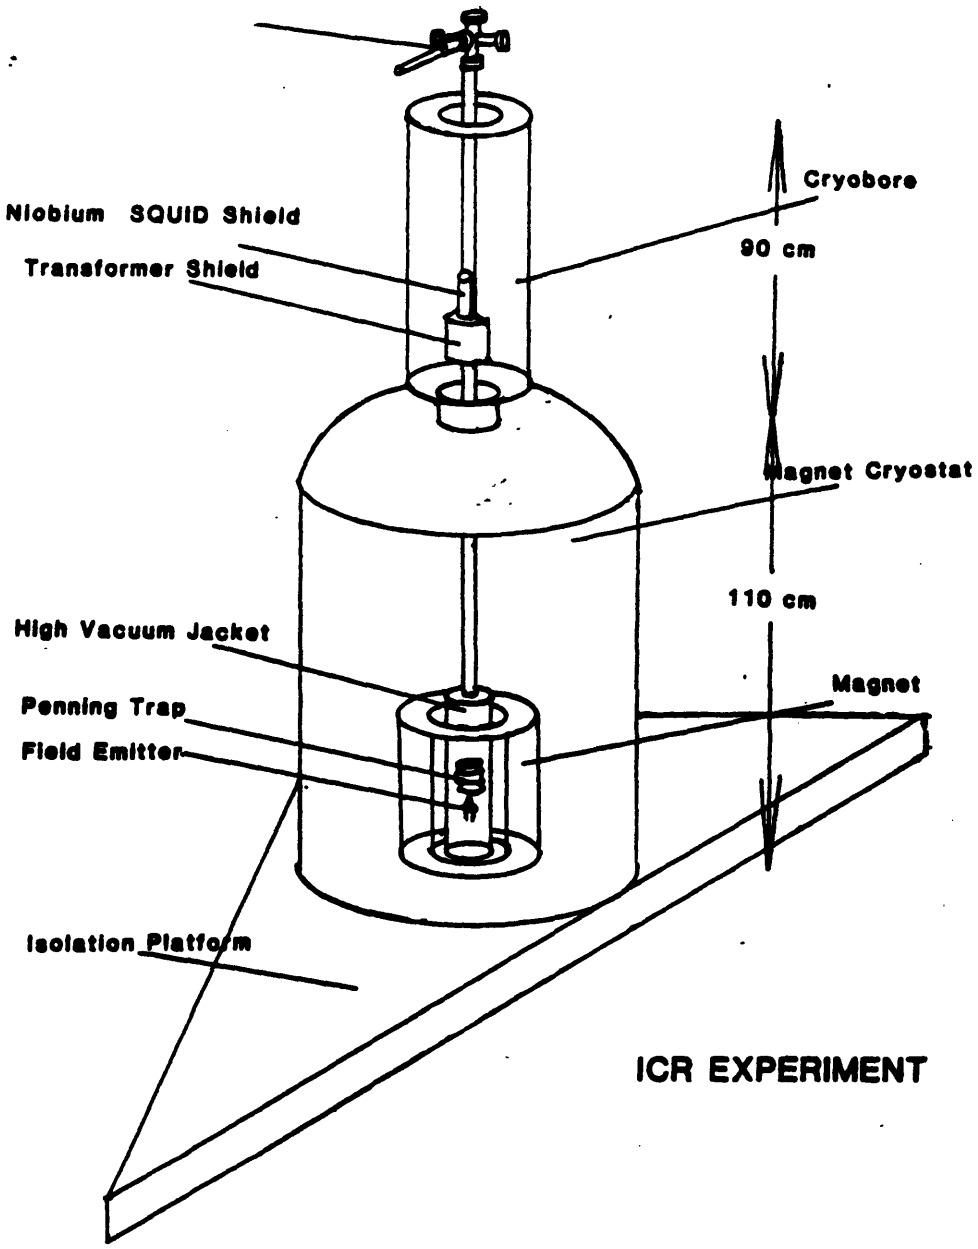

Figure IV.A.1. An overall sketch of the ion trapping apparatus, magnet, and cryogenic insert. ([FLA87], 3-9).

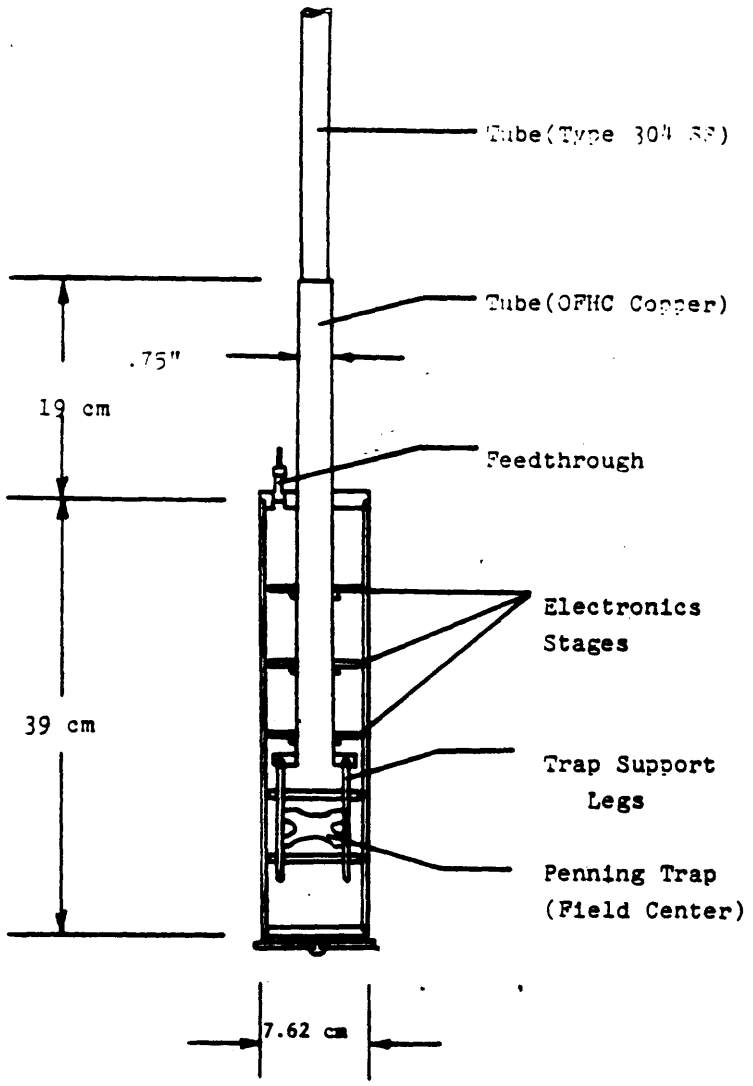

Figure IV.A.2. The Liquid Helium end of the probe. ([FLA87], 3-3).

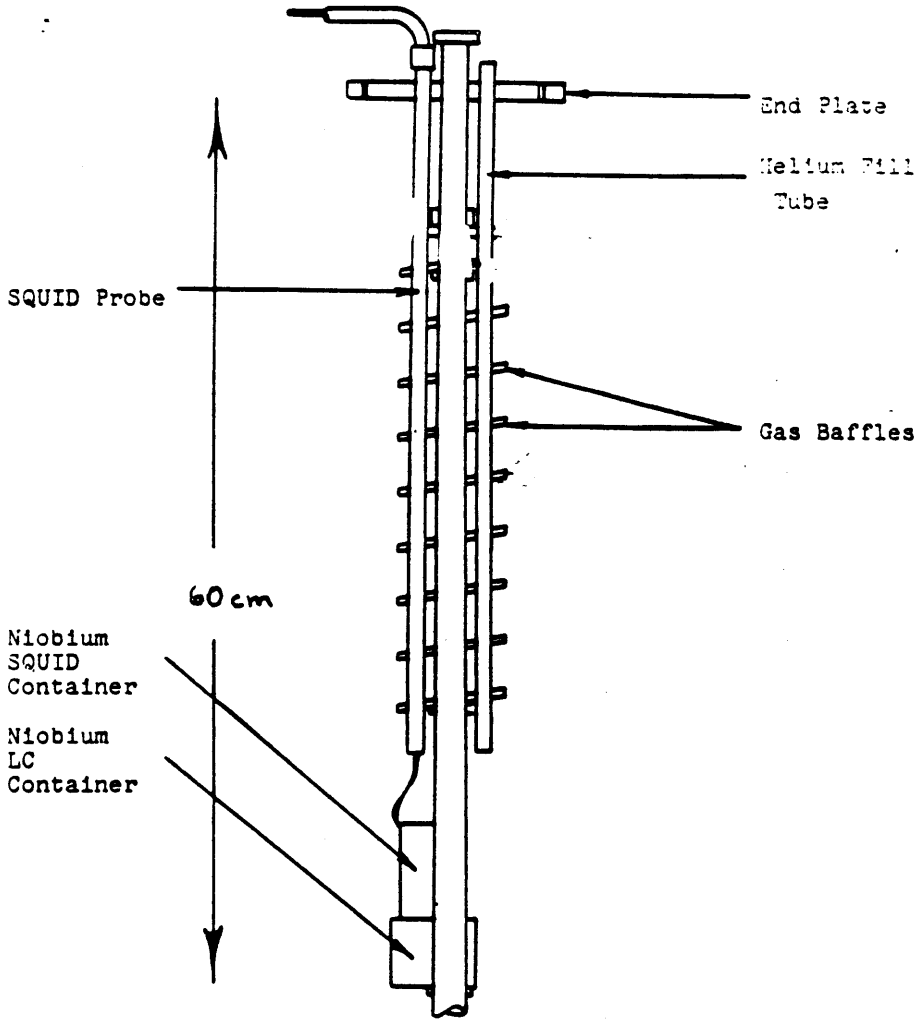

Figure IV.A.3. The upper part of the probe. ([FLA87], 3-4).

The top 40 cm of our vacuum tube is surrounded by a helix of G-10 disks, partially clad with aluminum foil. This helical baffle forces the cold helium gas to circulate around the experiment after the liquid boils, thus utilizing the enthalpy of the cold gas. The aluminum foil cut down radiation from room temperature.

We attached the cold SQUID sensor to its room-temperature pre-amplifier (the "RF-head") by a cryogenic, triax cable [note-spec] held fairly rigidly within small-gauge, stainless steel needle tubing filled with vacuum grease. This tube, in turn, was held within a 3/8" stainless steel tube by brass support disks placed at irregular intervals. (We hoped this irregular spacing would help minimize acoustic resonances within the probe.) The entire SQUID probe fit through a series of holes in the G-10 baffles.

Except for the connection between the trap and the detector, all the cables were run within S.S. needle tubing held inside a larger, 1/2" S.S tube. (See Figure IV.A.4) The stainless tubes end near the strong-field region of the magnet, and a 3/8" Copper tube continues down toward the trap. This tube simplified alignment with the trap, and reduced the likelihood that an errant wire could block the neutral atoms from entering the trap. The whole harness fit inside the main vacuum tube and was thermally connected to the it (and thus the Helium bath) at intervals with springy Be-Cu finger stock. This harness is an improved version of the earlier design in [FLA87].

The ion signal wires (Copper twisted-pair) came up though the OFHC can via a

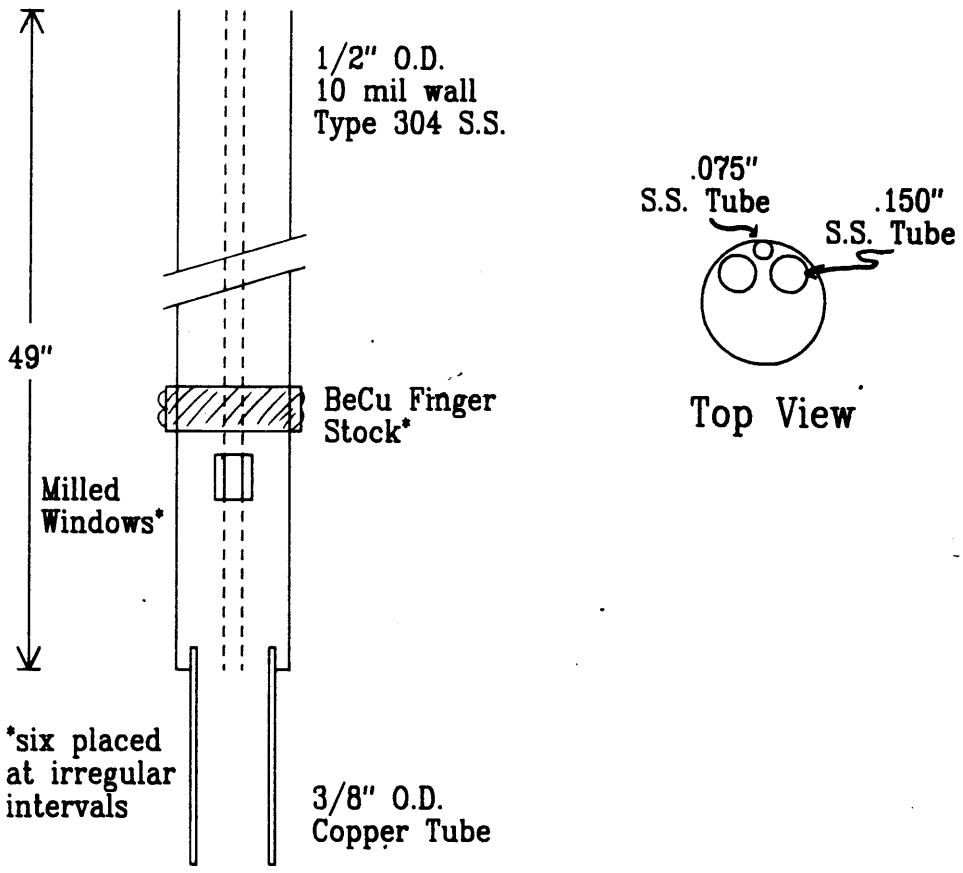

Figure IV.A.4. The new wiring harness.

high-vacuum feedthrough (mentioned in III.B, described in detail in FLA87), and passed directly though the liquid Helium bath to the detector. This twisted-pair was shielded within thin Copper refridgeration tubing. This shield was soldered at the detector end (to the same platform that supported the detector) and strain-relieved at the feedthrough end.

The trap biasing wires entered high-vacuum through another homemade feedthrough at the top of our probe. Since these wires were made from special, high-purity copper (an attempt to minimize thermoelectric effects), we kept them separate from the other wires. (We did not use these wires in the previous experiment.) The magnetron and cyclotron driving and coupling cables (more triax) entered though a BNC/mini-conflat feedthrough. The high-voltage cable for the field emitter entered through a similar MHV/mini-conflat flange. The axial drive cable and the DC guard ring wire entered through a commercial, multi-pin feedthrough.

To make ions, we set a negative potential difference between the trap ring and the endcaps, biased the field emitter to emit electrons, then, while electrons were whizzing up the field lines through the hole in the lower endcap, let in a small puff of gas at the top of the vacuum drift tube. Some fraction of the neutral atoms have trajectories in the correct direction (either initially or after a bounce or two) to enter the trap through a small, 1 mm hole in the upper endcap. Some positive ions then are made within the trap by electron impact ionization. Because the parent neutral had, at most, room temperature energy (25 meV), any atom ionized within the trap will

become bound there. (The ion's recoil energy from the ionization is negligible and the trap is at least 1 V deep, even for He.)

A few of the details on the old diagrams (Figures IV.A.2 and IV.A.3) no longer are correct. We have removed the optical baffle shown in Figure IV.A.2, originally designed to cut down room temperature radiation reaching the trap. This baffle made alignment especially difficult: several times it reduced the number of neutrals entering the trap so much that we could not make enough ions to detect. We decided that the small decrease in Helium boil-off was not worth this inconvenience. In fact, since the newer apparatus had about three times more cabling than the experiment described in [FLA87] (and thus a much larger heat load), we could not measure any appreciable increase in the liquid Helium boil-off rate without the baffle. The current helium boil-off rate is about 120 ml/hour.

In addition, the teflon spacer shown at the bottom of Figure IV.A.3 was removed[1] and two, more substantial spacers were machined and placed farther up on the stainless tube. We believe that these separated spacers force better alignment between our apparatus and the magnet bore, and therefore better alignment between the trap electrodes and the axis of the magnetic field.

IV.B Trap Modifications

1. Actually, it shattered and fell into the bore of the magnet

To shift the ion cloud, cool the magnetron motion and detect the cyclotron resonance, we made several modifications to the trap and its associated electronics. An additional high-purity Copper wire was brought down from room temperature with which we could set the lower endcap voltage independently from the ring voltage. This modification allowed us to move the equilibrium position of the ions within the trap, as described in Section II.B. In addition, we split the upper guard ring (by physically removing two small sections of the ring) to allow radial driving and cooling, described in Section VI.A and VI.B. See Figure IV.B.1 for a diagram of the completed trap and field-emitter.

Although it isn't immediately obvious, we can both couple and cool by splitting only one of the two guard rings. As discussed in Section VI.A, we *couple* the axial and radial modes with a quadrupole potential but *drive* the radial modes with a dipole potential. That is, the difference between direct driving and indirect coupling lies in the relative phase between the drives on upper and lower guard rings. (See Figure IV.B.2.) When we split only the upper guard ring, leaving the lower guard ring intact, we form a superposition of these two drives. (See Figure IV.B.3.) However, the driving and coupling resonances have quite different frequencies, separated in both the magnetron and cyclotron modes by ω_z. Thus when one drive (either the coupling or the direct drive) is resonant, the other must be so far from resonance that it can be neglected. Thus we simplified the wiring significantly by splitting only the upper ring.

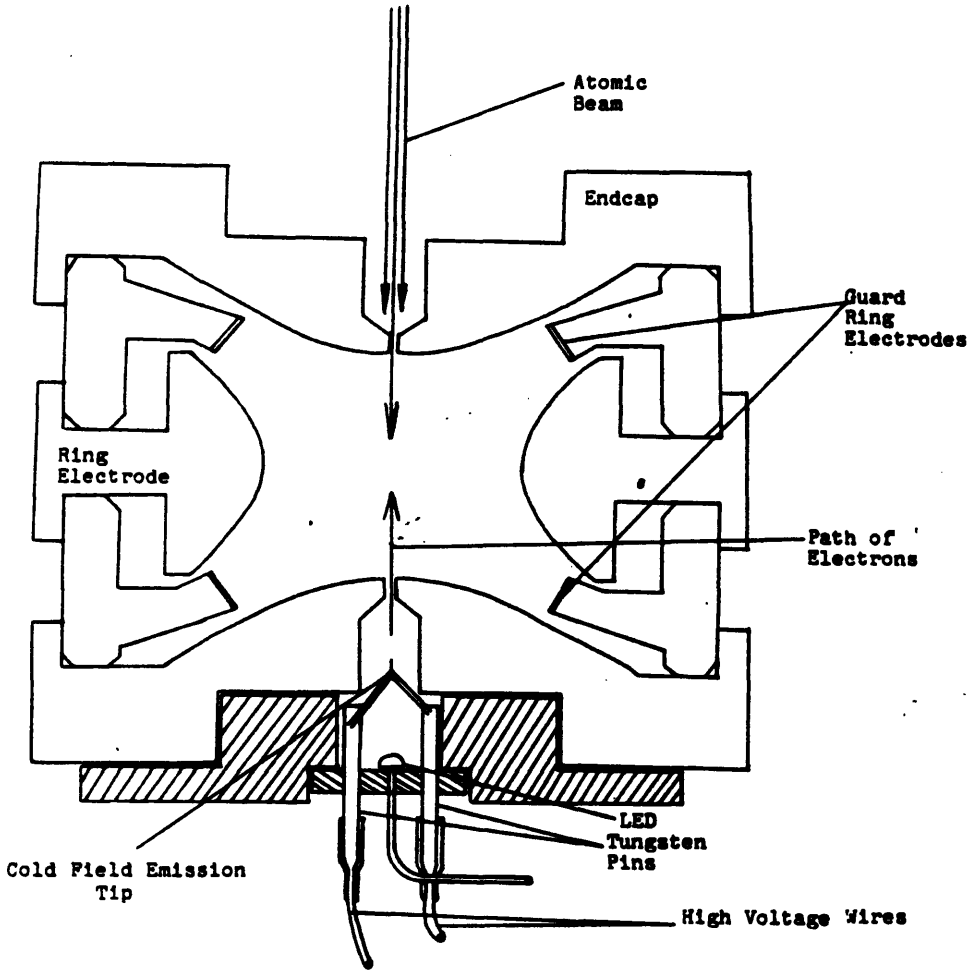

Figure IV.B.1. The Penning trap with field emitter mounted. The lower guard ring has been split, the gold surfaces painted with AquaDag, and the LED removed from below the FEP.

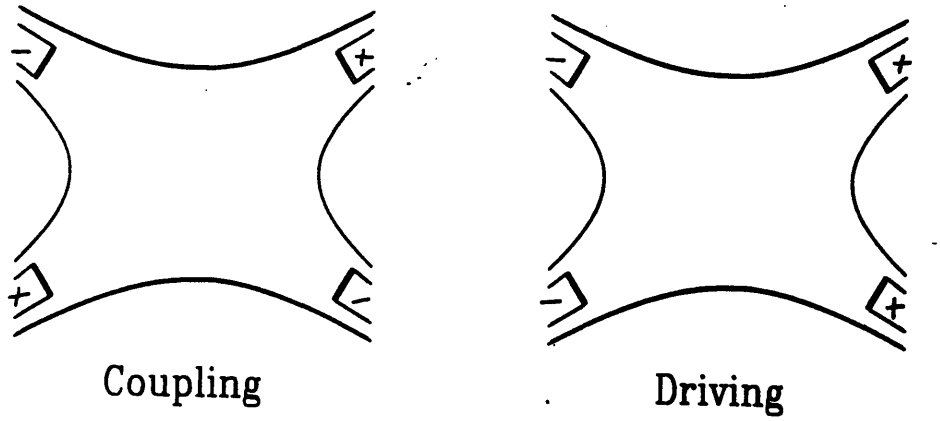

Figure IV.B.2. Phasing for coupling (left) and driving (right) the radial motions.

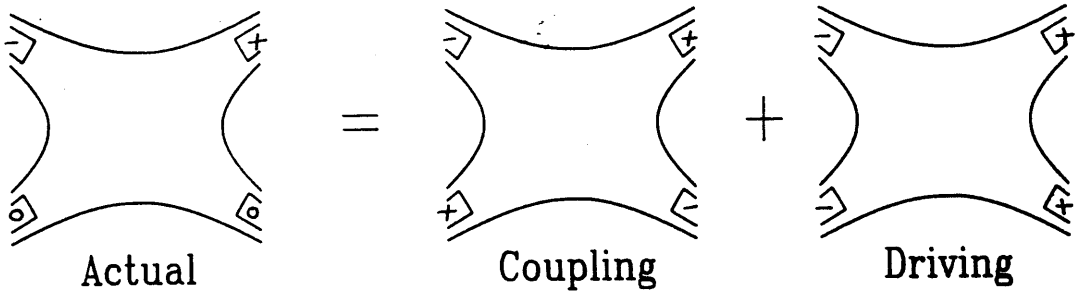

Figure IV.B.3. When only one ring is split, the guard ring drive gives rise to both driving and coupling.

The other major modification to the Penning trap was that we coated its electrode surfaces with Aqua-Dag, a carbon-colloid paint frequently used in spectroscopy experiments to blacken surfaces and suppress scattered light. We felt (the evidence will be presented Chapter VII; Section II.C.3 described the symptoms) that there was a significant charged patch on one or several surfaces of the trap. Although the trap was constructed from Copper for which one author [REF] found very little patch effects at 4 K, we had the trap plated with a thin gold layer. We discovered later that at low temperatures, such a surface can be particularly susceptible to patch effects. At about the same time, we heard that some groups [NRV87] studying Rydberg atoms (which are very sensitive to electric fields) found that Aqua-Dag effectively reduced patch-effect problems. Although their work was performed at room temperature, we felt it was worth trying at liquid Helium temperatures. We believe that the carbon paint significantly reduced built-up charges on the trap. (We give the evidence in Chapter VII.)

IV.C Trap Electronics

The added features—independent lower endcap, magnetron and cyclotron drives, axial/radial coupling—required additional low temperature electronics. These electronics, including some of the original filters mentioned in [FLA87] are shown schematically in Figure IV.C.1. The resistors are 1/4 W, metal film resistors which, unlike carbon film resistors, show good temperature stability and low 1/f noise at 4 K.

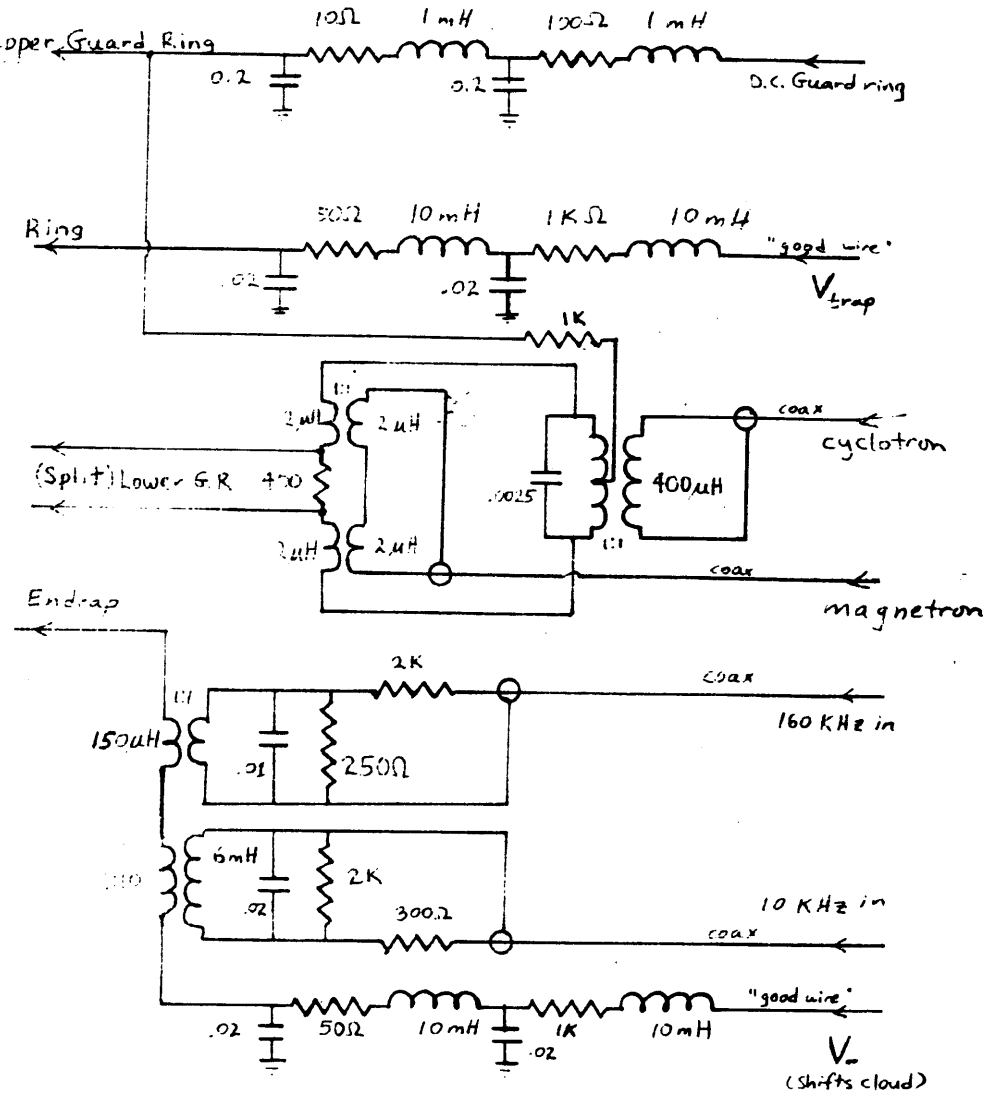

Figure IV.C.1. Schematic diagram for the cold electronics.

Most of the capacitors, originally mica capacitors, have been replaced by polystyrene capacitors, dunked in several coats of Torr Seal (a low vapor-pressure epoxy). These polystyrene capacitors are much smaller than the mica capacitors, and unlike the micas, do not have long nickel leads which must be removed before we could place them in the strong magnetic field. The inductors were wound on teflon forms (supplanting the old MACOR forms) with small-gauge, teflon-insulated Copper wire. To prevent them from moving, these forms were epoxied (again using TorrSeal) to copper holders.

The DC bias filters for the ring, lower endcap and guard ring all pass low frequency up to around 1 KHz, then cut-off dramatically at higher frequencies, as much as 80 dB per decade. We chose such a steep stop-band because a single ion will be significantly perturbed by small drives near its resonant frequency. For example, 1 μV of the right frequency surely could drive the ion strongly, perhaps even out of the trap. Even if the ion were not so responsive, the SQUID detector is quite sensitive to RF interference, and thus we tried to avoid all stray signals. We let the lower frequencies pass so that we could implement the two-drive detection scheme (described in Section V.D) with the local oscillator drive on the ring. For this scheme, we added a 1 KHz drive through the battery box at room temperature (Section IV.D) so that no additional low-temperature wires were required.

Notice from Figure IV.C.1 that we can also drive the lower endcap at 10 KHz. This extra feature was more important in the early days of the experiment when our

SQUID detector was so sensitive to RF interference that we could not always couple the 1 KHz though the ring. (This effect was never fully understood: the 1 KHz itself was not the likely cause of the interference; possibly the circuit coupling the 1 KHz at room temperature allowed higher frequency interference to propagate to the SQUID.)

IV.D New Voltage Source

We designed a better, more compact DC voltage source to bias the trap. This newer source was based on a particular precision reference integrated circuit, a Linear LT1021C, which provided extremely good temperature stability (5 ppm / °C) and very low noise (about 5 μV / pp). This residual noise, unfortunately is mostly 1/f and thus extra filtering does not lower its level. The resulting circuit fit inside a Copper box, 12" x 6" x 3", which we have put inside a second, acrylic box whose temperature can be controlled externally. Although noisier than a dry cell, the long-term (and as far as we could tell, even the short-term) stability of this source is far superior to the old Mercury cell source. A second function of this new "battery box" was to provide a convenient circuit to add a second, computer-controlled voltage (for scanning the ions' resonance) and the 1 KHz ring drive for two-frequency detection.

The voltage-source/voltage-adder circuit is shown in Figure IV.D.1. The trap voltage is set roughly by 10-turn, wire-wound potentiometer R1 which divides down the voltage given by the precision reference. Potentiometer R2 provides a fine adjustment to this voltage. The external AC and DC voltages are added into this

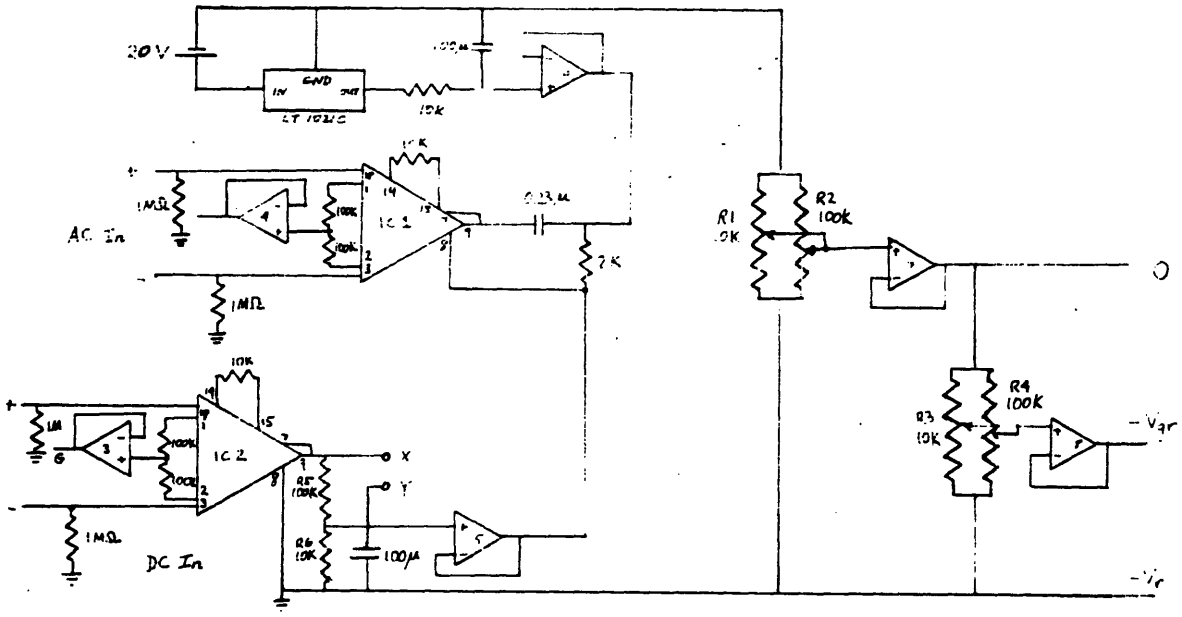

Figure IV.D.1. Schematic diagram for the "Battery Box," voltage reference and AC/DC offset circuit.

circuit using the offset capabilities of two AMP-01, ultra-low noise, high-stability instrumentation amplifiers, IC1 and IC2. Another feature of this circuit is that we can control externally a divide-by-eleven (through the resistive divider R5-R6) which lets us change between high- and low-resolution scanning. The divide-by-eleven, normally in, can be defeated by shorting points X and Y together. This short can be made manually or, as we do now, by an external relay driven by the computer.

The resulting voltage, buffered by an OP-27, ultra-low noise op amp, is sent down one good wire to the ring, shown on Figure IV.C.1. Within the voltage source, though, we divide the ring voltage by another pair of 10-turn pots, R3 and R4, buffer it by another OP-27, and use the resulting voltage to bias the guard ring. By dividing the trap voltage to produce the guard ring voltage, we keep their ratio the same. Thus, in the absence of patches, the trap should remain compensated during any voltage scan.

Though we could continue this buffering and dividing process one more time to set the lower endcap voltage, presently we change it so rarely that we set it externally by a simple 9 V transistor battery, divided down so that its small drift is unimportant. In addition, we have an auxiliary voltage, set by another battery, that we can activate with an external, computer-activated relay. In this way, we reproducibly can shift the ion cloud by remote control; a technique invaluable for the individual ion plateaus described in Chapter V.

The voltage source is powered by a set of rechargeable, 27.5 V gel cells whose outputs are regulated by a pair of ±15 volt regulators and a +20 volt regulator for the precision reference. Although there is some evidence that we cannot run for longer than 24 hours without recharging, as long as we do not need to make precision measurements, we can charge the batteries and, simultaneously keep the ions trapped.

There are two disadvantages of this voltage source. First, because we used "normal" 15 V op-amps, the total output voltage is limited to about 13.5 V, putting a limit on the detectable m/e at about 43 amu/esu. This limit must be overcome to study heavier molecules. Second, unlike batteries, our present silicon references have noise that limits the width of the axial resonance. This limit, for N_2^+ ions and the present detector, is just a bit below (a factor of three) a single ion's inductor-induced width, $\gamma_z = 0.3\ s^{-1}$. (This result improves at smaller masses. Because the bias is derived from the divided reference, the relative noise in the bias remains the same. However, the ion's width increase at lower mass, Equation II.A.3.7)

Both of these limitations can be overcome with appropriate technology. For example, Linear promises a plug-compatible reference with one-third the noise of the LT1021. Thus, when this noise becomes a problem, we will replace the chip.

IV.E Computer

We relied so much on our computer that it, too, must be considered part of the experimental apparatus. We used the computer to manipulate the fine-tuning of the

trap bias; sweep this bias in CW experiments; control the timing in pulse experiments (even more so in separated-oscillatory-field type measurements); it has formed an integral part of the narrow-band amplification of the ion signals; formed the central part of a feedback system to lock the ions to an external frequency source; and controlled the two, locked synthesizers necessary for the CW cyclotron resonance, among other things. In addition, the computer provided an integrated environment for simulating, recording, analyzing and presenting the data.

We used a DEC 11/73 microcomputer with various interface boards designed for real-time measurement and control: a Data Translation DT-2781, 16-channel A/D, 2-channel D/A board, configured to operate as 8 differential-ended inputs; an ADAC 1601GPT general purpose timing board to synchronize the experiment; a DEC DRV11-C Parallel Input/Output board provided 16 independent control bits for setting relays, starting timers, etc.; an MLSI GPIB interface board to communicate with our synthesizers (two HP3325A, DC-60 MHz multi-function synthesizers) and, optionally, a digital voltmeter (Keithley 197, 5 1/2 digit DVM).

The 11/73 ran under a UNIX-type operating system (called "VENIX") which created a powerful system both for developing programs (simulations, analysis and presentation) and for controlling the experiment (writing flexible, real-time control code). Clearly this thesis is not the place to present reams of programs any more than it is the place to show the schematic of every electronics box we slapped together for a special purpose. However, several programs (and external hardware we used to

cooperate with them) were sufficiently important and general that they are worth a few paragraphs here.

Low Pass Filters

Narrowing the bandwidth of a detected signal is one the most common methods used to improve signal-to-noise. Since measured white noise has an rms amplitude that depends on the bandwidth of the detector, Δf, like $\Delta f^{-1/2}$, we see that decreasing Δf can decrease the average noise without attenuating the signal. When we are sampling the signal with the computer (a process discussed in Section III.B), it would be convenient to allow filtering of the signal with software since the signal already has been measured and stored within the computer. There is a whole sub-specialty in Electrical Engineering devoted to the design of digital filters, and we took advantage of their work [OPS75, DSP79] to design several particularly simple low-pass filters.

Because we used a general-purpose computer (rather than a machine designed specially for signal processing) and because we did not want to become experts in either multiplication-optimization or round-off error-propagation, we chose to implement infinite impulse-response (IIR) filters rather than finite impulse-response (FIR) filters. In general, FIR filters allow better phase control across their frequency response (because they have no poles, only zeros, they have constant phase); however they also require far more multiplications than an IIR filter that implements the same amplitude response. With an IIR filter, we needed so few arithmetic operations that we could use floating-point multiplication and by-pass many round-off problems.

The first type of filter we used was a (very) low pass filter designed with phase-sensitive detection in mind. Although there is extensive lore about such filters (see, for example, VOW86), much of the lore concerns squeezing the last few percent of efficiency out of such filters. Thus for these long time-constant filters (tau between 1/4 and 25 seconds), we used the digital equivalent of a simple RC-filter, designed using the conservative bilinear transformation procedure. This procedure takes the poles of a continuous filter and maps them into the z-domain (the discrete-time space analogous to frequency space in continuous-time, and thus appropriate for digital filters). We then implemented the resulting z-domain pole using a direct form. [OPS75] The c program text for this filter is given in Figure IV.E.1. The subroutine filt_init(), called before the filter is used, sets the proper constants within the filter and initializes the storage. Every time a data point is read from the A/D board, it is filtered using filt_go() which returns the filtered value

The second type of low-pass filter we used served a different purpose. As we discussed in Section III.B, undersampling a signal with higher frequency components can lead to aliasing; that is, signals sampled at less than twice their frequency will appear at lower frequencies. Noise and coherent signals alike will alias, therefore merely sampling at a low frequency does not perform any frequency narrowing: it just aliases all the higher frequency noise into a narrower spectrum.

We often sampled signals at around 1 KHz, and used an analog filter (discussed later in this section) designed to minimize aliasing by sharply attenuating any signals

```
#define PI      3.14159265358979
#define C(n,i)  sav[n+n+n+i]

static double   sav[3],fb,gain;

/*
 * filt_init() must be called first to initialize the filter for
 * the user-specified time-constant.
 *              "freq" is the sampling frequency
 *              "tc"   is the time constant
 */
filt_init()
{
        double  omega;
        omega = 2. * tan( (double) 1. / ( 2. * freq * tc ) );
        gain = omega / ( 2. + omega );
        fb = ( 2. - omega ) / ( 2. + omega );
        sav[0] = sav[1] = sav[2] = 0.;
}

/*
 * filt_go() is called each time a new point is to be filtered.
 * It returns the filtered value.
 */
float filt_go()
float   x;
{
        double  in;

        in = x * gain;

        C(0,1) = C(0,0);
        C(0,0) = fb * C(0,1) + in;
        in = C(0,1) + C(0,0);

        return( (float) in );
}
```

Figure IV.E.1. Code for quick low-pass filter.

above about 500 Hz. However, sometimes we wanted to look carefully at lower frequency signals, say a 50 Hz wide part of the spectrum. In order to avoid aliasing we designed a "down-sampling" filter that mimicked sampling at much lower filters. This filter first performs a sharp cut-off of signals above about 1/4 of the sampling frequency. Since this step filters out the higher frequency information, we can throw away three of every four filtered points: they are redundant. This filtering-and-discarding process, called down-sampling, effectively lowers the sampling frequency without the risk of aliasing.

The c code for the down-sampling filter is shown in Figure IV.E.2. The sharpness of the cut-off required cascading several stages together. We used a "Direct Form-II" [OPS75] implementation that minimized the round-off errors in successive cascades. The theoretical and measured response of this filter is shown in Figure IV.E3. The type of filter (Chebychev) gives a sharp cutoff at the expense of ripple in the pass band. However, for precision application we can use Figure IV.E.3 to compensate for the ripple when analyzing the amplitude of a response at any known frequency. The 11/73 is fast enough so that we can cascade two of these down-sampling filters together to make a 1/16 down-sampler, too.

For both of these filters, we made no determined attempt to optimize either for speed or for the application. In much the same way that we might solder together a few resistors and capacitors (for the first filter), or spend a day or two piecing together some op-amps and other components (for the anti-aliasing filter), these filters barely

```
#define C(n,i)  sav[n+n+n+i]

static double   sav[6];

/*
 * filt_init() must be called first to initialize the filter
 */
filt_init()
{
        register        i;

        for (i=0; i<6; i++)
                sav[i] = 0;
}

/*
 * filt_go() is called each time a new point is to be filtered.
 * It returns the filtered value.
 *      This filter lets the user safely down-sample by a factor
 *      of four.  As can be seen from the constants within the
 *      subroutines, it is quite inflexible.
 *      In fact, this subroutine was generated by another program.
 */
float filt_go()
float   x;
{
        double  in;

        in = x * 0.001836;              /* note = proton-electron mass ratio */

/* Iteration 1 */
        C(0,2) = C(0,1);
        C(0,1) = C(0,0);
        C(0,0) = -0.848200 * C(0,2) + 1.499600 * C(0.1) + in;
        in = C(0,2) + C(0,1) + C(0,1) + C(0,0);

/* Iteration 2 */
        C(1,2) = C(1,1);
        C(1,1) = C(1,0);
        C(1,0) = -0.649300 * C(1,2) + 1.554800 * C(1,1) + in;
        in = C(1,2) + C(1,1) + C(1,1) + C(1,0);

        return( (float) in );
}
```

Figure IV.E.2. Code for 4-pole Chebychev, down-sampling filter.

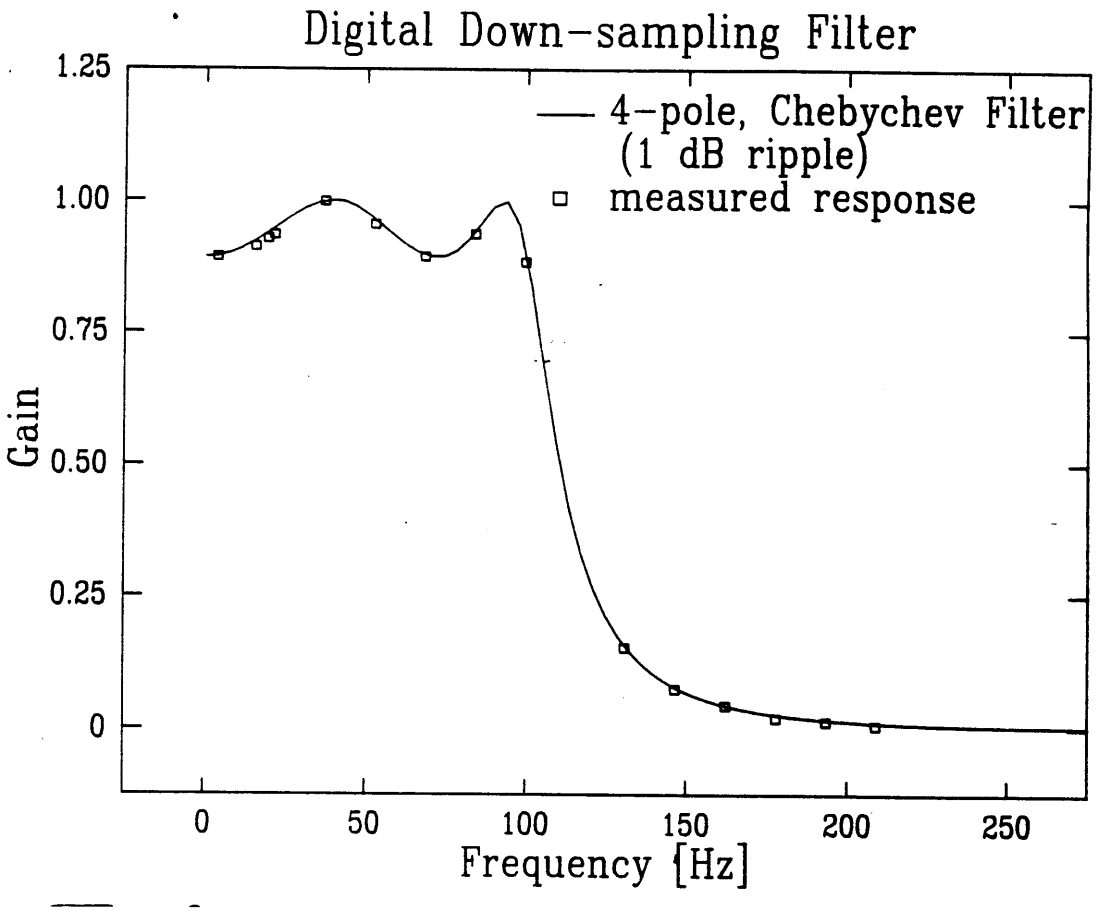

Figure IV.E.3. Response of the digital filter shown in Figure IV.E.2.

scratch the surface of what we might be able to do, should the need arise.

Phase-Sensitive Detector

As discussed in Section V.C and VI.B, we often measure the CW response of the ions by driving them and detecting their response at a known frequency. In this case, we'd like to use a phase-sensitive amplifier to narrow the bandwidth of the detector and thus improve the signal-to-noise of the result. Early in the experiment, we discovered that we could not use a commercial lock-in amplifier because the extra RF noise it generated made it impossible to use the SQUID. We therefore implemented a simple computerized lock-in using an external frequency generator, a mixer, an anti-aliasing low-pass filter and the computer. (See Figure IV.E.4) The frequency generator, mixer, and low-pass filter generate a low-frequency image of the signal, typically around 250 Hz. Starting with a signal at ω_{signal}, when we multiply it with a mixing signal at ω_{mix}, we get two strong harmonics out at $\omega_{signal} \pm \omega_{mix}$. The low-pass filter throws out the high frequency signal at $\omega_{signal} + \omega_{mix}$, passing the difference signal to the computer.

Although any reasonable low-pass filter would eliminate the summed harmonic, to avoid aliasing, we used a 4-pole Chebychev filter, described in [HOH80], implemented as shown in Figure IV.E.5. Like the digital anti-aliasing filter described above, it has ripples in the pass band that must be remembered in precision work. The predicted and measured responses are shown in Figure IV.E.6. (Notice how much better we can predict a priori the response of digital filter, Figure IV.E.3!)

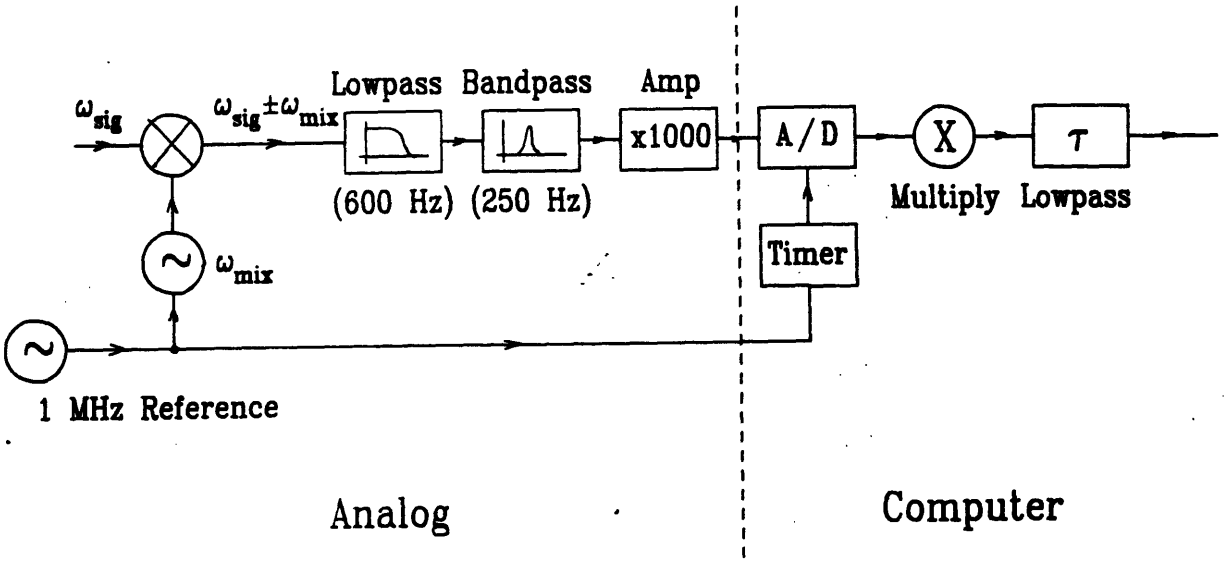

Figure IV.E.4. Diagram of phase-sensitive detector.

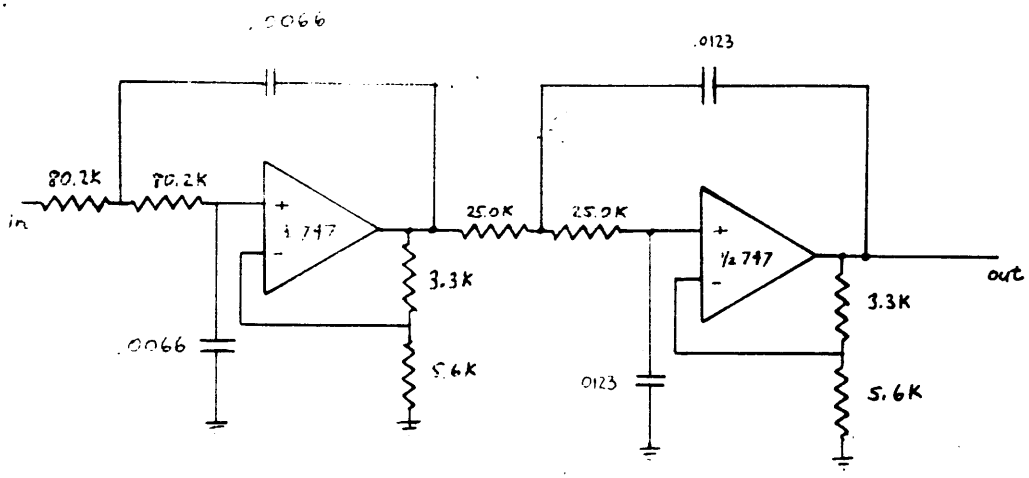

Figure IV.E.5. Schematic for two op-amp, four-pole Chebychev, 0.5 dB ripple anti-aliasing filter.

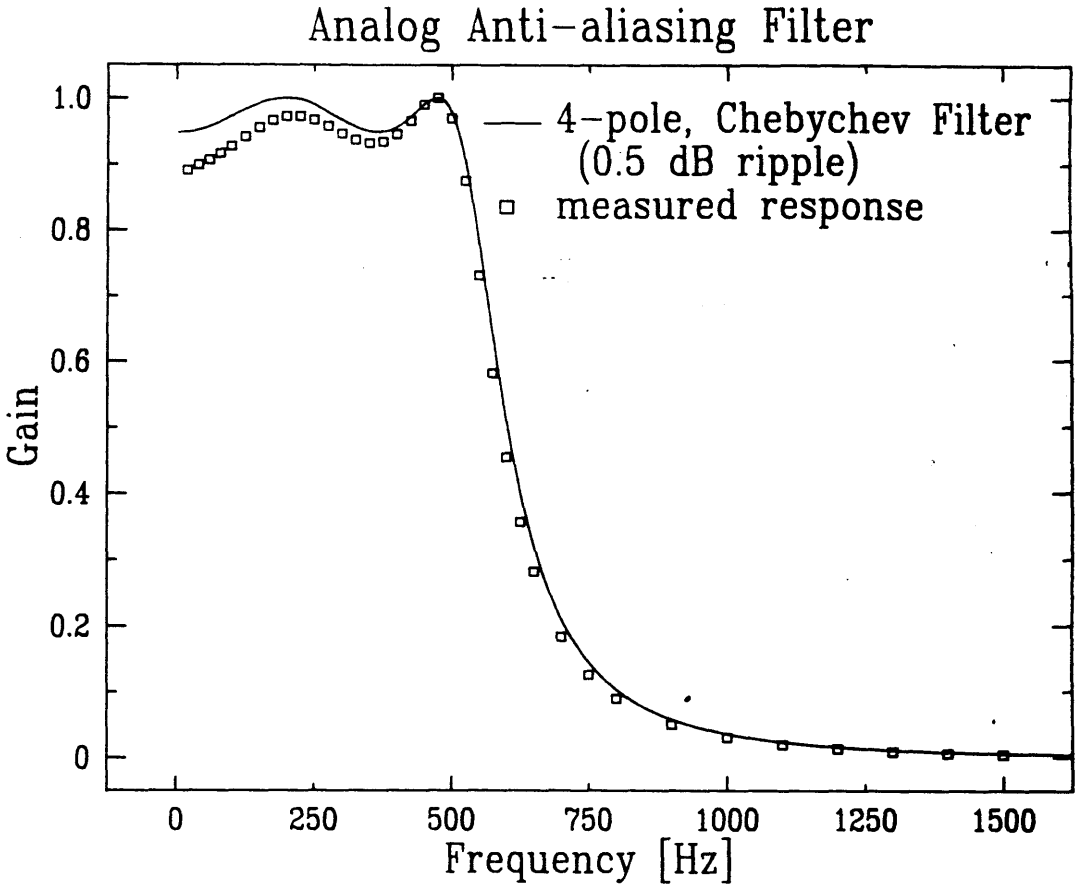

Figure IV.E.6. Response of the analog filter shown in Figure IV.E.6.

The computer then performs another mixing step at $\omega_{signal} - \omega_{mix}$ by multiplying by sine and cosine functions at the appropriate frequency, set by the user. The results are passed to the first low-pass algorithm described above, resulting in a phase-sensitive signal at DC.

We do the mixing in two steps—first with analogue mixers, then digitally—for two reasons. First, we avoid the 1/f noise of the amplifiers, both the Evans amplifiers on the schematic, Figure IV.E.4, and also those on the A/D board. Second, and more importantly, we avoid generating explicitly the ions' resonant frequency. In the two-drive scheme, the ion mixes together two different frequencies, and is resonant when the sum of these frequencies matches the ions' natural frequency. To detect, we again mix with two different frequencies that sum to the ions' response frequency. At no time do we explicitly generate the ions' response frequency. Thus, in theory at least, we never see a sum signal except when the ions themselves generate it.

In practice, however, there can be two sources of "false" signal. First, the two driving signals, ω_{LO} and ω_d from Section V.C, can mix together a little bit, either in the output stage of the frequency generators or at the SQUID itself. (Each has non-linear elements, and thus can generate harmonics.) Second, because we often use ω_{LO} to clock the A/D's (otherwise we run out of frequency generators), the small fraction of ω_d which is not attenuated by the low-pass filter will appear aliased at the same frequency that the ion should respond. However, neither of these "signals" will tune as we change the trap potential and thus these red herrings rarely cause great trouble.

Coherent Transient Digitizer

For pulsed experiments, we cannot use phase-sensitive detection because we do not know ahead of time the true frequency of the response. For example, when we pulse the ion near its resonant frequency (as described in Section V.A.3), we expect the ion and the detector to ring at their natural frequencies. Since, after all, we are trying to measure these frequencies, we cannot know the phase of the response. In that case, then, we must improve the signal-to-noise by averaging together several independent runs. Even when we do not know what the frequency will be, we can pulse the ions in ways that will guarantee (in the absence of drifts) that the phase will be identical from run to run. The signals will add coherently, noise will add in quadrature, and thus we get the usual $\sqrt{N}$ improvement in signal-to-noise.

For example, in the two-drive pulse scheme, Section III.E, we synchronize the two drive frequencies with a third frequency so that both drive frequencies are multiples of this synchronizing frequency. In addition, we synchronize the start of the pulse itself with this third frequency. The resulting ion response will have the same phase each pulse, and therefore we can add several consecutive responses together to improve the signal-to-noise of the result. This process is shown diagrammatically in Figure IV.E.7. The pulse-synchronizer, a simple TTL device lifted from [HOH80], is given in Figure IV.E.8.

The pulse-synchronizer acts like a switch. It takes a series of clock pulses

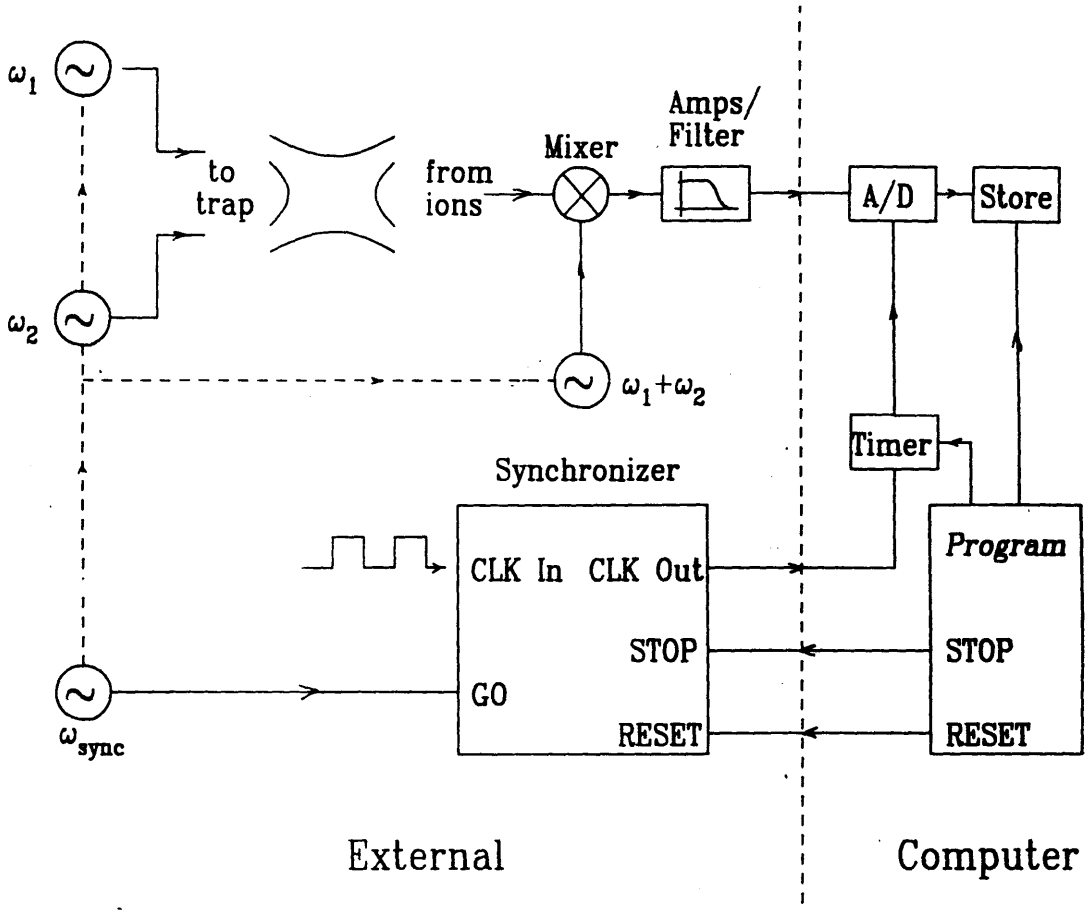

Figure IV.E.7. Diagram of coherent transient averager.

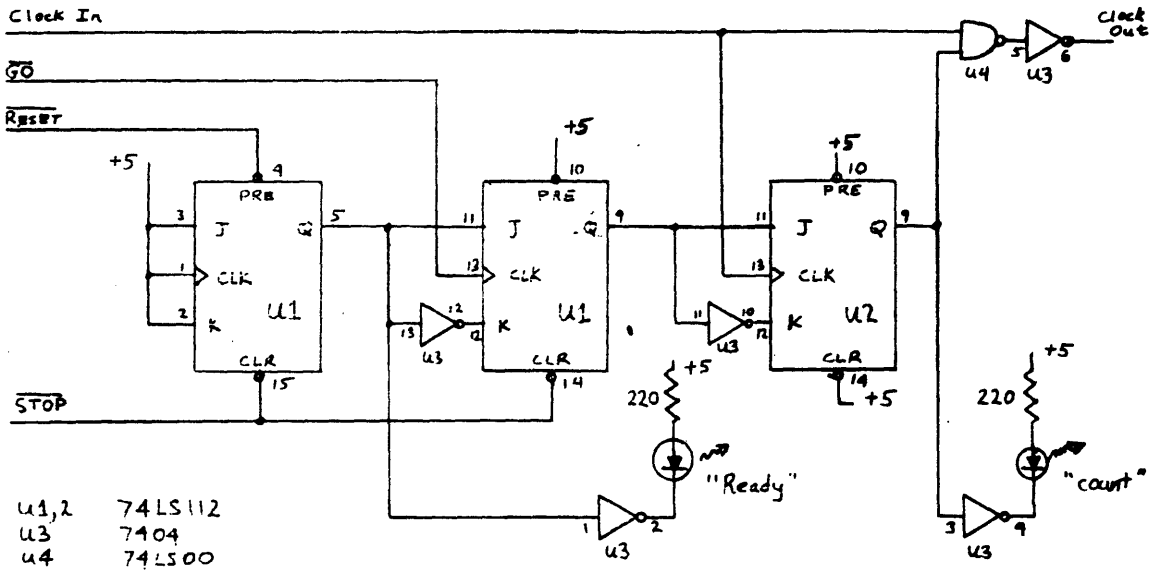

Figure IV.E.8. Schematic diagram for pulse-synchronizer.

(usually 1 MHz pulses) and lets them through them under certain conditions. When the synchronizer is stopped, nothing passes. The pulse-synchronizer waits until it gets a "RESET" pulse from the computer. After this pulse, the next state-change on the synchronizing input "GO" lets through the 1 MHz clock pulses. These clock pulses are used as the timing source for the general purpose timer board in the computer. That timer, in turn, controls when the A/D's begin each analog-to-digital conversion. For pulse detection, we use the same analog mixer and low-pass filter described for the lock-in. The coherent averaging then takes place in the computer, set up by the pulse-synchronizer.

When the desired amount of data has been taken, the computer sends a "STOP" pulse that the pulse-synchronizer uses to cut off the 1 MHz clock. The pulse-synchronizer now ignores all input until it receives another RESET pulse, and the cycle begins anew. In this way, we have synchronized the data-taking with the ions' excitation. The largest phase error we can incur should be due to one missed cycle of the 1 MHz clock out of the signal the computer measures; that is, the mixed, low-passed signal. Since this signal always is less than 250 Hz, the phase error is negligible. We then fourier transform the summed responses (See Section V.A) and can determine the frequency, amplitude, and phase of the result.

IV.F Procedure

In this section, I describe the general procedure we used to prepare the apparatus

for detecting ions. When warm, the apparatus hangs from a shelf in a lab one floor above the magnet. Assume the worst: that we have taken the penning trap off the rest of the probe, disassembled the trap, and taken the field-emitter (FEP) off the bottom endcap. (In fact, the last time we completely warmed the apparatus—February, 1988— this was the state of affairs.)

We remounted the FEP, carefully fitting it into the recess in the lower endcap. (Figure IV.B.1 shows the location) We reassembled the trap in the proper order: the tight tolerances we required could only be met for a given orientation and assembly order for each of the pieces. The FEP holder was then adjusted so that the place where the FEP is spot-welded could be seen through the upper endcap. (The FEP itself is far too small to be seen end-on.)

We then placed the trap on the three threaded copper rods on the probe, Figure IV.A.2, using three threaded washers to place the trap at the location which should be magnetic field center. Three springs above the trap force it together, and must be "sprung" to assure that the trap holds together. We then used the threaded washers to level the trap's upper surface, first with respect to the top plate two meters above, then with respect to the perpendicular of the S.S vacuum drift tube. Because this drift tube is slightly bent, these two methods gave slightly different results, and thus we tried to compromise between the two. Finally, we sighted the FEP holder from the *top* of the drift tube. We have mounted a window in a mini-conflat right at the top, and using a prism and an alignment telescope, we checked that we could see all the way down the

drift tube, out it's end, beyond the top of the trap (which we sighted first, because its gold surface was illuminated easily), then, by lighting the trap from below, the FEP itself. We were convinced that if we could see the FEP, neutral atoms would have a path by which they could meet the beam of electrons within the trap. When we cannot sight the top of the trap, we must adjust a series of set screws drilled into the copper part of the drift tube (within the OFHC vacuum can, of course) that push at different angles against the bottom of the wire harness, Figure IV.A.4.

We then tested, at room temperature, all of the "cold" electronics. Since it usually takes several days to cycle the experiment from warm to cold and back to warm again, we tested as much of the circuitry as we could: it seemed that every time we were lazy at this step, we paid for that sloppiness in wasted time. We made sure that the DC lines were continuous to the trap and had the appropriate resistances. We checked that the drive lines had the proper frequency responses. (Since the components do not vary much when cooled to 4 K, these measurements also serve as absolute calibrations so that we can predict the amplitude of the response when we drive the ions.) Currently, the damage we find is usually limited to exploded resistors. (In olden days, we sealed the can at much higher temperatures, and often accidentally unsoldered connections. We switched to a higher melting-point solder[2] for connecting

2. "Stay-brite"

the electrical components to minimize this risk.) At this point, we also checked that the SQUID excitation circuit would tune at the appropriate frequency (19 MHz).

After checking the electronics, we sealed the OFHC can onto the probe with Indium metal and StayClean flux (a general purpose, acid flux). To avoid overheating the experiment or turning the trap into an oven, we follow a fairly rigid procedure to seal the can which is quite different from the one described in [FLA87]. Unfortunately, we did not consider a good sealing surface when we designed the experiment, and thus must rely on the roundness of the can and the tightness of its fit to the thick Copper shelf (See Figure IV.A.2) to provide a good seal. (In fact, we had to have a new can made when the first two became sufficiently out-of-round to seal without gobbling up $50 worth of Indium.)

We cleaned the tinned surfaces with Acetone and then carefully fitted the can up around the trap and electronics and against the Copper shelf. We held the can in place on a platform mounted on a ring stand. Using Propane torches we heated the Copper block until the flux wetted the surfaces. The can, which has much less thermal mass, was heated indirectly by contact with the Copper block and by the edge of the torch flame. Once parts of the block reach the melting temperature of Indium, (the top of the can always reached this temperature first) some parts sealed themselves with the Indium left from the previous seal. Where there was insufficient Indium, we added small amounts of new metal and small amounts of flux. We believe the upper part of the block will be just above the melting point of Indium, while the lower parts just

below; this temperature differential created an effective sealing surface.

Usually, there were several small regions along the interface across which it was obvious that the seal would leak. Often, Indium flowed badly at these points because there was some junk just below the seal or the can was particularly out-of-round at these points. Heating those spots directly and re-fluxing apparently brings these points just above the Indium flow temperature leaving neighboring regions just below; sealing up these small imperfections. We let the seal cool, cleaned it with warm water, then Acetone. We leak-tested the clean joint (and the room-temperature conflats, too) and, when necessary, repaired any pin-holes in the seal, again by local heating.

Once convinced we had a seal, we baked the apparatus gently at about 90 C for 4-6 hours, to drive off solvents and water, especially from the trap surfaces and the VYCOR pieces. (The VYCOR, a porous glass, sits loose at the bottom of the can, providing extra 4 K surface area to pump away He and H when the experiment is cold.)

After the experiment cooled from the bake-out, we painted the exposed Indium surface (and some of the brazes) with a low-temperature, polystyrene-based cement[3] which we think minimized the risk of any additional leaks when the apparatus becomes surrounded by liquid Helium. We then put about half a liter of liquid

3. Kurt J. Lesker Co. Leak Sealant

Nitrogen into the (pre-cooled) cryogenic insert in the magnet's bore. (This amount of Nitrogen should cool the experiment from room temperature to 77 K, leaving no liquid Nitrogen left in the bore). We slowly lowered the experiment through a 16" hole in the concrete floor directly into the bore of the magnet. (The magnet lives on the floor below.) The boiling liquid Nitrogen cools the apparatus to near 77 K. We monitored the temperature with a thermocouple near the bottom of the apparatus, and added liquid Nitrogen if needed.

We let the apparatus rest at 77 K for at least twelve hours to let the long time-constant parts of the dewar and the apparatus equilibrate. After determining that there was no liquid Nitrogen left in the dewar (or blowing any out if there were), we transferred liquid Helium (about 20 l) into the bore, passing the Helium through a long extension on the transfer tube to the bottom of the bore. For the first 24 hours, the Helium boil-off rate was always quite high, reflecting, we believe, the long time it takes the dewar's radiation shields to cool. Although we could trap ions at this stage, the extra noise from the boil-off (and the lower Q for the detector, as mentioned in the last chapter) usually force us to wait a day.

After the boil-off had settled down, we determined the resonant frequency of the detector (it varied by about 500 Hz from cool-down to cool-down), measured its Q, checked for stability and, using past ion data, set the trap bias voltage to bring the desired ions into resonance with the detector. At this point, we set the guard ring voltage to its last good setting.

Using the procedure outlined at the beginning of this chapter, we then tried to make a relatively large cloud of ions (say about 40) and tried to see their noise (the damping of their initial energy), or, failing that, a pulsed or CW response. Since the trap dimensions may have changed a small amount from cool-down to cool-down (or there can be irreproducible patches), we sometimes have to search a bit to find the ions.

Having found the ions, we then tuned the trap (for example, see results in Section VII.A) by making swept, two-drive CW resonances as symmetric as possible with these larger clouds. As discussed in Section II.B, smaller clouds become increasingly sensitive to anharmonicity, and thus small mistunings of the guard ring can obscure such clouds entirely. Once the guard rings have been set with moderate clouds, we make smaller clouds and use their increased sensitivity to improve further the tuning.

In the worst instances, the guard rings *began* quite mistuned, and thus we saw very little. Unfortunately, mis-set guard rings mimic other problems (like contamination by other ions or a small Helium leak), and thus we must spend considerable time searching carefully through parameter space (both ring and guard ring settings) to try and identify the difficulty. When all seemed amiss, we usually found that making large clouds (~100 ions) and driving them right after they were made very strongly off-resonance would inevitably give us *some* signal. When this signal decayed over a period of several minutes and could not be found again, the

problem was usually contamination by Helium ions. However, when that signal showed some stability, then sweeping the ions (see the next chapter) and shifting the guard rings in roughly 1% steps usually produced good results. However, it appeared that when the guard rings were mistuned, impurity ions problems, especially drift in the frequency of the "good" ions (see Section II.C.2) became particularly bad. Therefore, we found that making a new batch of ions between each guard ring change produced better results sooner.

Once the ion resonance has been found and the guard rings roughly tuned, the experiment remains quite stable. That is, the uncertainty in trapping potentials comes from the initial cool-down, and the settings, once found, need only minor adjustment, at least for the several month periods for which we have kept the experiment cold. Once we have found the ions' resonance, we can begin to try the more interesting experiments.

CHAPTER V

DETECTING THE AXIAL MOTION

In the next two chapters, we will discuss in detail detecting and cooling the ions in the trap. In this chapter, we focus on the axial motion, deferring the radial motions to Chapter VI. Within this chapter we will introduce theory to explain several different schemes to excite and measure axial motion, discuss their advantages and drawbacks, and present measurements made with each scheme.

We discuss in Section V.A driving the ions with a single, oscillating electric field. Though this scheme is very simple to implement, the interaction of the ions, the detector, and the trap capacitance make the resulting detector current somewhat complicated to interpret. We demonstrate that the ions can be treated like a series LC-circuit. Using this model, we then determine the detector and ion currents in both the time and frequency domains.

In Section V.B, we describe an experiment that used pulsed, single-drive excitations to measure the natural width several-ion clouds. In that section, we also present data demonstrating several of the novel conclusions from the theory in Section V.A.

In the rest of the chapter, we discuss two-drive schemes. For these methods, in addition to the drive discussed in the previous sections, we also modulate the trapping potential. In Section V.C, we show that this method allows us to induce ion motion without exciting the detector with the drives. Although the equations of motion

become time-varying, and thus not amenable to normal fourier transform analysis, we present an approximate Green function that, for these purposes, lets us treat the system simply.

We present results using this two-drive scheme in Section V.D. The first result, unfortunately buried in the middle of this chapter, is our ion counting data, showing steps due to single ions within the trap. We also show in Section V.D, practical examples of several of the ideas from Chapter II. We show anharmonic resonances and show that the guard rings have be used to make the trap more harmonic. We demonstrate that moving the ions off the coil produced narrower resonances. We also describe methods to lock the axial resonance to an external frequency source.

This chapter concludes with a discussion of two-drive pulsing, Section V.E. This technique has many of the advantages of single-drive pulsing without the additional complication of the detector/ion coupling. We show (theory and experiment) that we can improve the signal-to-noise in this scheme by putting the ions off the detector's resonance. These final data also show some drawbacks, both for two-drive pulsing and for discrete fourier transforming.

V.A One Drive Techniques

The current which trapped ions induce in the endcaps in response to an applied voltage has the same form as a series *lc*-circuit. Viewing the problem in this manner lets us understand more easily the motion of a trapped ion hooked up to a detector and

drive. The response of ions to an external drive yields several surprises, especially when the ions are driven by a single oscillating voltage, simultaneously resonant with both the ions and the coil. This equivalent circuit approach was introduced in Wineland and Dehmelt's ion calorimetry paper [WID75], and is expanded below.

V.A.1 Equivalent Circuit for the Ions

We begin with the results from Section II.A.2. The ions can be affected by two different electric forces. (See Fig V.A.1.1) First, the quadrupole trapping potential yields a restoring force, $F_z = -m\omega_z^2 z$. In addition, when there is an axially anti-symmetric potential V across the endcaps, the ions will experience a force, $\vec{F}_{appl} = e\,\vec{E} = -\frac{eB_1}{2z_0} V\,\hat{z}$, where B_1, as explained in section II.B, is a constant that expresses, to first order, the difference between a Penning trap and a parallel plate capacitor.

Combining these two forces, we have for the z motion:

$$\ddot{z} + \omega_z^2 z = -\frac{eB_1}{2mz_0} V \qquad \text{(V.A.1.1)}$$

In addition, as discussed previously and summarized in (II.A.3.4), the ion's motion induces a current:

$$I = -\frac{e\,B_1}{2z_0}\,\dot{z} \qquad \text{(V.A.1.2)}$$

Now consider a configuration in which the trapped ion has been replaced by a series *lc*-circuit, Figure V.A.1.2. Ignoring the trap capacitance, we can solve for the current,

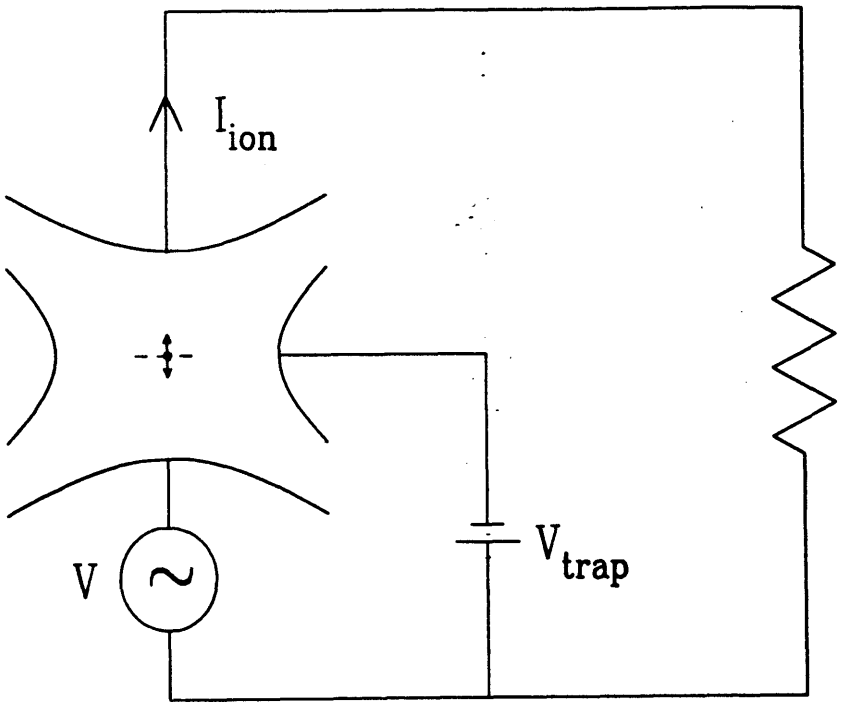

Figure V.A.1.1. The ion held within the trap. V_{trap} biases the trap to hold ions. V, oscillating near the trapped ions' resonance frequency, drives the ions. As it oscillates, an ion induces a current, I_{ion}.

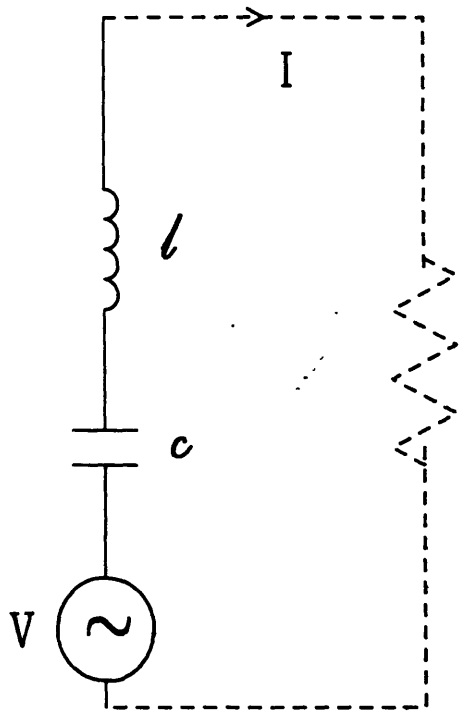

Figure V.A.1.2. Electrical equivalent circuit for the trapped ion in Figure V.A.1.1. The ion has been replaced by a series *lc* -circuit.

I, in terms of the driving voltage, V:

$$\ddot{I} + \frac{1}{lc} I = \frac{1}{l} \dot{V} \tag{V.A.1.3}$$

Taking the time derivative of (V.A.1.1) and using (V.A.1.2) to replace $\dot{z}$ with I:

$$\ddot{I} + \omega_z^2 I = \frac{e^2 B_1^2}{4mz_0^2} V \tag{V.A.1.4}$$

Comparing (V.A.1.3) to (V.A.1.4) we see the strong similarity. Exploiting this similarity, we can write:

$$\begin{cases} l = \dfrac{4mz_0^2}{e^2 B_1^2} \\ \\ c = \dfrac{1}{\omega_z^2 l} \end{cases} \tag{V.A.1.5}$$

For one N_2^+ ion in our trap, $l = 4.1 \times 10^8$ H, $c = 2.4 \times 10^{-21}$ F. Although such values might appear to be too extreme to have any interactions with the detector or the trap, we shall see, in fact, that unlike the case of trapped electrons discussed in [WID75], these extreme values will have a profound effect on detection.

For more than one ion in the trap, as discussed in Section II.A.2, the induced current increases. Indeed, for N ions, the current (I in Equation V.A.1.2) will be N times larger. Thus, when we make the substitution into Equation V.A.1.2, we find:

$$\begin{cases} l \rightarrow l/N \\ c \rightarrow N\,c \end{cases} \qquad \text{(V.A.1.6)}$$

That is, the equivalent inductance decreases inversely with the number of ions and the capacitance increases linearly. The resonant frequency stays the same. We also can relate γ_z, the damping introduced by detection, to the ions' inductance, l. In Section II.A.3 we showed

$$\gamma_z \equiv N \cdot \frac{1}{\mathrm{m}} \left[\frac{eB_1}{2z_0}\right]^2 \mathrm{Re}Z \qquad \text{(V.A.1.7)}$$

But from (V.A.1.5) and (V.A.1.6) above, we see readily that:

$$\gamma_z = \frac{\mathrm{Re}\,Z}{l} \qquad \text{(V.A.1.8)}$$

This relation seems the most compact way of expressing the fact that the ion damping comes directly and naturally from the detector.

The strength of this approach is that it will let us treat the ions, the trap, and the detector on equal footing, using all the tools (and short cuts) of linear circuit theory. We always can recover the ion's motion using (V.A.1.2) to determine the ion's velocity from the current. I use this approach in the next section to determine both the motion of the ions and the detected signal when a single, resonant voltage (CW or pulsed) is used to drive the system.

V.A.2 Single Drive Green Function

In this section, I will determine the response of the ions and the detector to a single drive. This problem is important for two reasons. First, we frequently measure the axial motion with this kind of drive (usually with a pulse), and, as we shall see, this response can contain significant information. Second, the structure of the solution, and in particular, the Green function for the problem, is very similar to the structure for other problems in the experiment. For example, magnetron cooling and two different kinds of cyclotron resonance have similar structures. Therefore, the discussion of the solution to this problem will be useful for the other problems, too.

Because the theory is central to our various detection schemes, this section is rather long. In it, I extend the equivalent circuit method of the last section to include the trap capacitance and our resonant detector. After making some general observations about ion/detector coupling and the effects of trap capacitance, I determine the detected current, first in the frequency domain, then in the time domain. I then proceed with a similar analysis for the ion's motion. The analysis will point out several important effects which have not been mentioned in the literature, including the strong dependence of the driven motion on the trap capacitance and a weaker dependence on the number of trapped ions.

Let me first explain the full equivalent circuit. In the experiment, we frequently detect the axial motion by driving the ion from the lower endcap and detecting the current induced in the upper endcap. (see Figure V.A.2.1) Following the results of the previous section, the trapped ion can be replaced by a series *lc*-circuit in parallel with

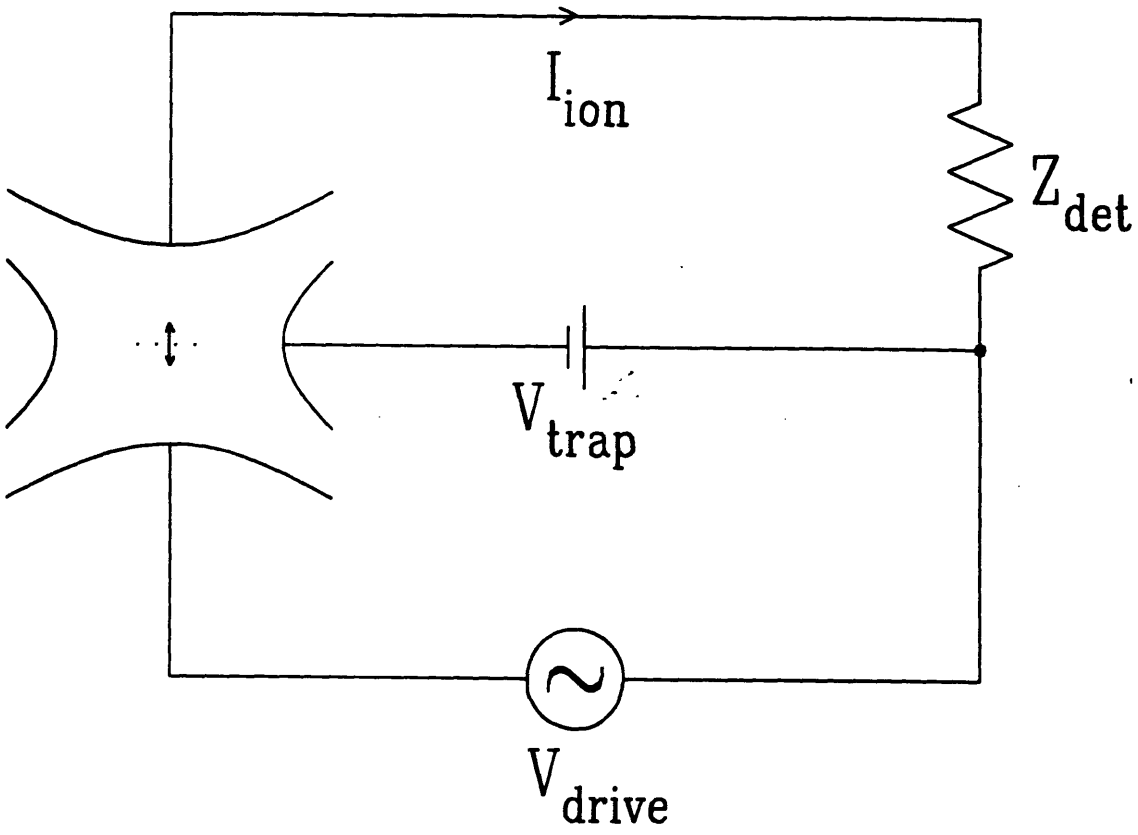

Figure V.A.2.1. Physical set-up for driving and detecting trapped ions.

the trap capacitance, C_{trap}. (C_{trap} actually is the sum of the physical capacitance of the trap plus the parasitic capacitance of the drive and detector cables.) In addition, I replace the resistive detector of Figure II.A.3.2 with a more accurate model. This detector is a parallel-resonant, LC-circuit, with some finite loss, which I represent as a parallel resistor, R. (Since we will concentrate on effects near the detector's resonance, it is unimportant whether the loss is modeled as a large, parallel resistor or a much smaller resistor, in series with the detector capacitance. The actual mechanism of the loss is inconsequential near resonance.) Finally, since we measure the current through the detector inductor, L, I will concentrate here on determining that current. The complete equivalent circuit is shown in Figure V.A.2.2.

Before grinding through the algebra, let me make a few guiding observations. This problem is essentially that of two coupled harmonic oscillators; one damped, one undamped. In such problems, there are usually two distinctly different regimes: a "weakly-coupled" regime, corresponding to motion that resembles the independent oscillations of the constituent oscillators; and a "strongly-coupled" regime, corresponding to combined motions: symmetric and anti-symmetric sums.

We can estimate the transition condition between these two regimes for the ion/detector system. I argued before (Equation V.A.1.8) that the width of the ion, because of the detector, is $\gamma_{ion} = \frac{\mathrm{Re}\,Z}{l}$. Right on the tuned circuit resonance, $\mathrm{Re}\,Z$ is at its peak value, $\mathrm{Re}\,Z = R = \omega_0 LQ$. However, that argument assumed that the

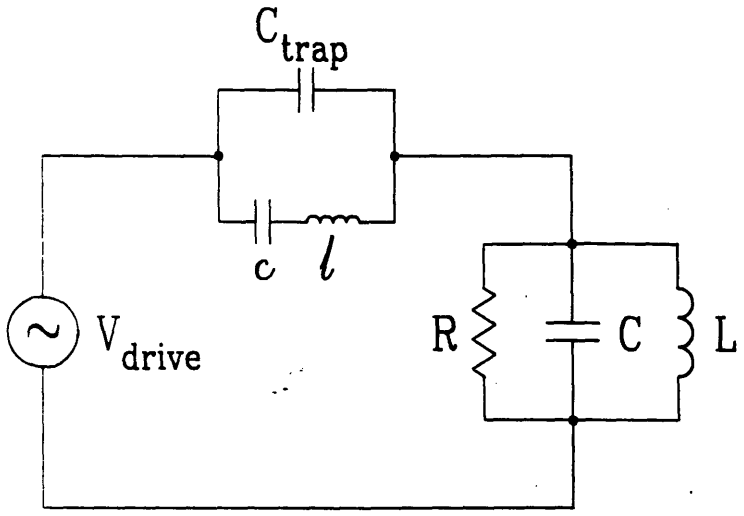

Figure V.A.2.2. Equivalent circuit for situation in Figure V.A.2.1. The trapped ion has been replaced by a series *lc*-circuit, in parallel with the trap capacitance. The detector has been replaced by a parallel *LC*-circuit with finite dissipation.

induced width, γ_{ion} is narrower than the coil width, $\gamma_{coil} = \frac{\omega_0}{Q}$. When this is not the case, the phase argument used in section II.A.3 to derive γ_{ion} is not valid. We can use $\gamma_{ion} \ll \gamma_{coil}$ as the weak-coupling limit. In terms of the equivalent circuit components, this limit yields $l \ll Q^2 L$. Since we operate with $Q > 10^4$, the apparently huge value of l for the ions isn't necessarily that large after all. Indeed, our experiment presently crosses into the strongly-coupled regime at about 30 N_2^+ ions, and, since l scales with the mass of the ions, this limit will be only three $^3H^+$. (These values should be compared to $\approx 10^4$ e^- in the Washington experiments. [WID75])

Let me also make two important observations about the trap capacitance. Neglecting the ions for a moment, what happens at the resonance of the LC-circuit? At that frequency, $\omega_{LC} = (LC)^{-1/2}$, the detector has impedance R (the contributions from the inductor and capacitor cancel), and thus the voltage across the detector is given by:

$$V_{det} = V_{in} \frac{i\,\omega_{LC} R C_{trap}}{1 + i\,\omega_{LC} R C_{trap}} \qquad \text{(V.A.2.1)}$$

Here, $\omega_{LC} R C_{trap} = Q \frac{C_{trap}}{C}$. Since Q is so large, it is likely that $\omega_{LC} R C_{trap} \gg 1$. (Currently, this limit requires only $C_{trap} \gg 6\times10^{-15}$ F, a condition certainly met.) Thus $V_{det} \approx V_{in}$, and there is *no* voltage drop across the trap: drives near ω_{LC} should not affect the ions because there will be no electric field to act on them.

But ω_{LC} is *not* the resonance frequency of the detector. Since the trap capacitance is in parallel with the tuned circuit capacitance, the total resonance frequency shifts down to $\omega_0 = [L(C + C_{trap})]^{-1/2}$, as we find when we calculate, in general, the voltage drop across the detector:

$$V_{det} = V_{in} \frac{-\omega^2 L C_{trap}}{1 - \omega^2 L [C + C_{trap}] + i\omega L/R} \qquad \text{(V.A.2.2)}$$

Thus, on the detector's resonance, ω_0:

$$V_{det}(\omega_0) = V_{in} \cdot i\omega_0 R C_{trap} = V_{in} \cdot Q \frac{C_{trap}}{C} \qquad \text{(V.A.2.3)}$$

As discussed above, $Q\frac{C_{trap}}{C}$ can be much larger than one. Therefore, at the detector's resonance, the drive the ions see can be many times larger than the drive applied on the lower endcap.

In passing, note that we can use Equation V.A.2.2 to determine the Q of the detector. Since the current through the inductor is $V_{det} / i\omega L$, we get:

$$I_{det} = \frac{i\omega C_{trap}}{1 - \omega^2 L [C + C_{trap}] + i\omega L/R} V_{in} \qquad \text{(V.A.2.4)}$$

That is, the voltage drive acts like an current source, $V_{in} i\omega C_{trap}$, and this current is amplified by the resonance of the detector. Thus, despite the fact that the absolute impedance of the trap capacitance is much smaller than the impedance of the tuned circuit on resonance, as long as $C_{trap} \ll C$ this small impedance does not "ruin" the high-Q resonance of the circuit.

Bearing these points in mind, we can calculate the voltage drop across the detector for a given drive voltage. Dividing that voltage by the impedance of the inductor yields the current through the inductor, I_L:

$$I_L = \frac{i\omega(c + C_{trap})(1-\omega^2 l c_p)}{(1-\omega^2 l c)(1-\omega^2 L[C + C_{trap}] + i\omega L/R) - \omega^2 L c} V_{drive} \qquad \text{(V.A.2.5)}$$

where $c_p \equiv \frac{C_{trap}\, c}{C_{trap} + c}$ is the effective capacitance of the trap and the ion in series. Since $C_{trap} \gg c$ in most cases, c_p is indistinguishable from c. The denominator of (V.A.2.5) has two resonant terms, corresponding to the detector and the (undamped) ions. The third term, $\omega^2 L c$, represents the coupling between the oscillators, and hence depends on both the detector, through L, and the ions, through c.

We can rewrite (V.A.2.5) in a more tractable form in the narrow-resonance approximation, which assumes that ω is near the natural ion resonance frequency, ω_z, and the detector resonance frequency, ω_0. We can also simplify by noting that $C_{trap} \gg c$. Thus, near resonance, we have:

$$I_L = \frac{i\omega_0^2 C_{trap}}{2} \frac{\Delta_z}{\Delta_z(\Delta_0 - i\gamma_0/2) - \gamma_0\gamma_z/4} V_{drive} \qquad \text{(V.A.2.6)}$$

where the "0" subscripts denote coil parameters:

$$\omega_0^2 = \frac{1}{L(C + C_{trap})} \qquad \text{(V.A.2.7)a}$$

$$\gamma_0 = \frac{\omega_0^2 L}{R} \qquad \text{(V.A.2.7)b}$$

$$\Delta_0 = \omega - \omega_0 \qquad \text{(V.A.2.7)c}$$

and the "z" subscripts denote ion parameters:

$$\omega_z^2 = \frac{1}{lc} \qquad \text{(V.A.2.7)a}$$

$$\gamma_z = \frac{R}{l} \qquad \text{(V.A.2.7)b}$$

$$\Delta_z = \omega - \omega_z \qquad \text{(V.A.2.7)c}$$

The parameter γ_z would be the damping of the ions if they were coupled to a purely resistive detector as described above. This parameter is the true ion damping only when $\gamma_z \ll \gamma_0$, as we shall see below.

Equation V.A.2.6 summarizes the entire linear dynamics of the ion/detector system. In fact, the admittance (the ratio of current to voltage) is the fourier transform of the Green function G(t - t′):

$$I(t) = \int dt' G(t - t') V(t') \qquad \text{(V.A.2.8)}$$

The poles of the admittance identify the resonances of the system. The real part of a pole corresponds to the resonant frequency, and the imaginary part corresponds to half the damping. In Equation V.A.2.6, the poles are simply the zeroes of the denominator:

$$\omega_{1,2} = \frac{1}{2}\left\{\omega_z + \omega_0 + \frac{i\gamma_0}{2} \pm \sqrt{(\omega_z - \omega_0 - i\gamma_0/2)^2 + \gamma_0\gamma_z}\right\} \qquad \text{(V.A.2.9)}$$

The relative size of the two terms under the radical determines the coupling regime.

When the first term dominates (large ion-coil detuning or very small ion damping), we call the system weakly-coupled. In this case, we can expand the square root in (V.A.2.9) to lowest order:

$$\omega_1 \approx \omega_z + \frac{1}{4}\frac{\gamma_0\gamma_z}{\omega_z - \omega_0 - i\gamma_0/2} \qquad \text{(V.A.2.10a)}$$

$$\omega_2 \approx \omega_0 + \frac{i\gamma_0}{2} - \frac{1}{4}\frac{\gamma_0\gamma_z}{\omega_z - \omega_0 - i\gamma_0/2} \qquad \text{(V.A.2.10b)}$$

Here, ω_1 clearly corresponds to the ions and ω_2 to the coil. There is a small frequency shift in both resonances, $\pm\mathrm{Re}\left[\frac{1}{4}\frac{\gamma_0\gamma_z}{\omega_z-\omega_0-i\gamma_0/2}\right]$ which is almost always inconsequential. (For example, the axial frequency shifts at most a few parts in 10^8 per N_2^+ ion.) More importantly, the ions are given a width (and the coil is narrowed by):

$$\gamma_{eff} = 2\,\mathrm{Im}\left[\frac{1}{4}\frac{\gamma_0\gamma_z}{\omega_z-\omega_0-i\gamma_0/2}\right] = \gamma_z\frac{\gamma_0^2/4}{(\omega_z-\omega_0)^2+\gamma_0^2/4} \qquad \text{(V.A.2.11)}$$

Thus, when the ions are near the coil, $|\omega_z-\omega_0| \ll \gamma_0/2$, the effective ion damping becomes $\gamma_{eff} \approx \gamma_z$, as discussed in Section II.A.3. On the other hand, when the ions are moved off the coil, $|\omega_z-\omega_0| \gg \gamma_0/2$, their damping will be reduced by the square of the detuning. This reduction exactly corresponds to the reduced real impedance of the tuned circuit off its resonance. Thus the ion's damping always comes from the real part of the detector impedance.

We can summarize the weakly-coupled limit by rewriting the Green function:

$$I_L \sim \frac{1}{\Delta_0 - i\gamma_0'/2} \frac{\Delta_z}{\Delta_z - i\gamma_{eff}/2} V_{drive} \qquad \text{(V.A.2.12)}$$

The first term in this expression is the normal, lorentzian-like response of the coil, with a slightly reduced damping. In the second term, on the other hand, the zero in the numerator dominates at the ion's resonance: on their resonance, the ions "short out" the coil and no current flows in the detector. The total frequency response for the weakly-coupled case is shown in Figure V.A.2.3.

However, when enough ions are in the trap (presently, $\gtrsim 30\ N_2^+$ ions), the second term in the radical of Equation V.A.2.14 dominates. In this case, neglecting the first term entirely yields:

$$\omega_{1,2} = \frac{\omega_0 + \omega_z}{2} \pm \frac{\sqrt{\gamma_0\gamma_z}}{2} + \frac{i\gamma_0}{4} \qquad \text{(V.A.2.13)}$$

These modes represent combined motions of the ion and detector, with eigenfrequencies split above and below the average of the uncoupled modes. The amount of that splitting increases slowly with γ_z. These two "strongly-coupled" modes have the same damping, $\gamma_0/2$. Thus the ions can *never* be given a damping from the detector larger than half the natural damping of the detector. (We will find a similar limit when we discuss cooling the magnetron and cyclotron modes.) We can rewrite the admittance in this limit:

$$I_L \sim \frac{\Delta_z}{(\Delta_1 - i\gamma_0/4)(\Delta_2 - i\gamma_0/4)} V_{drive} \qquad \text{(V.A.2.14)}$$

where $\Delta_{1,2} = \omega - \omega_{1,2}$. Thus the strongly-coupled response resembles two identical

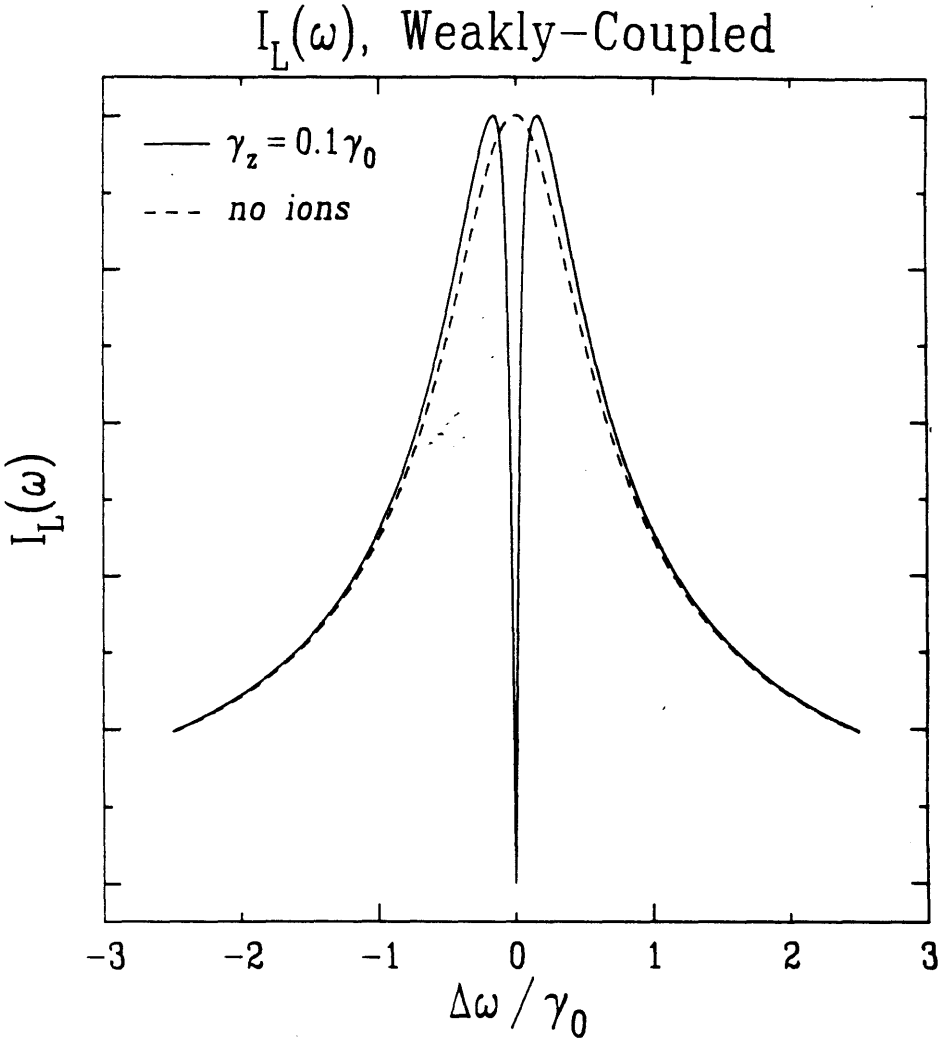

Figure V.A.2.3.a The spectrum of the detected current in the weakly-coupled case. The dashed line shows the lorentzian response of the detector when no ions are in the trap. The solid line shows the ions "shorting-out" the detector on resonance.

detectors, split by the presence of the ions, each with half the damping of the real detector. Notice, however, that on resonance the detector is still shorted out by the presence of the ions. (see Figure V.A.2.3b)

We can use the admittance above to determine the time-domain response of the coil. Since the admittance represents the fourier transform of the impulse response, we can transform back to get the response in the time domain. When we do the inverse fourier transform, terms like $\frac{1}{\omega-\bar{\omega}-i\Gamma/2}$ become, in the time domain, $e^{-\Gamma t/2}e^{-i\bar{\omega}t}$. We shall write frequency responses (V.A.2.13 and V.A.2.15) as the sum of two such terms, the time-domain response therefore is the sum of two dying sinusoids, and thus may destructively or constructively interfere.

For both the weakly- and tightly-coupled regimes, then, we seek to factor the coil admittance:

$$G(\omega) = \frac{(\omega-\omega_z)}{(\omega-\omega_1)(\omega-\omega_2)} \qquad \text{(V.A.2.15)}$$

into two terms:

$$G(\omega) = \frac{A}{\omega-\omega_1} + \frac{B}{\omega-\omega_2} \qquad \text{(V.A.2.16)}$$

The algebra is straight forward, and yields:

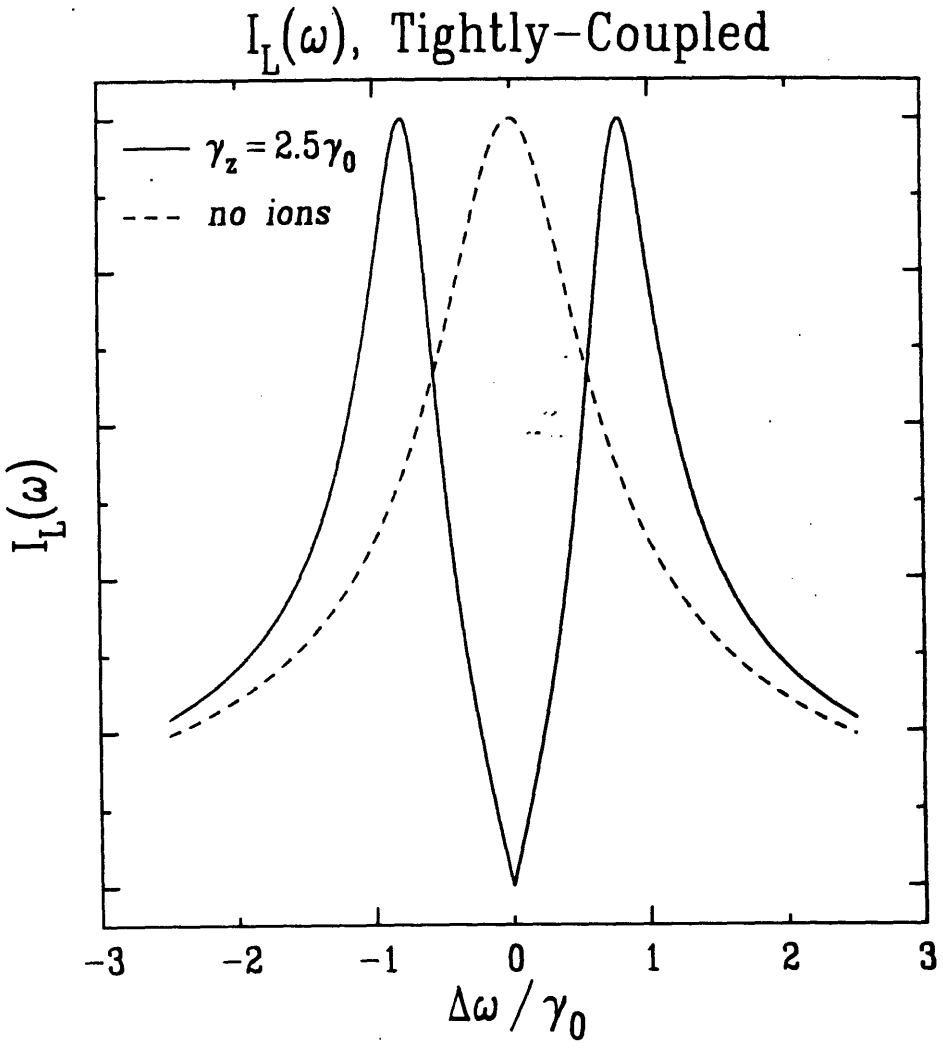

Figure V.A.2.3.b The spectrum of the detected current in the tightly-coupled case.

$$A = \frac{\omega_z - \omega_1}{\omega_2 - \omega_1} \qquad \text{(V.A.2.17)a}$$

$$B = \frac{\omega_2 - \omega_z}{\omega_2 - \omega_1} \qquad \text{(V.A.2.18)b}$$

Although the relative size of A and B is set by the relative proximity of $\omega_{1,2}$ to the uncoupled ion frequency, the absolute size of both will be set by the difference between the two coupled modes, $\omega_2 - \omega_1$.

Placing the ions right on the coil, $\omega_z = \omega_0$, we now analyze the two coupling regimes separately. In the weakly-coupled case, $\omega_1 = \omega_0 + i\gamma_{eff}/2$, $\omega_2 = \omega_0 + i\gamma_0/2 - i\gamma_{eff}/2$, and the time domain response will be:

$$I(t) \sim e^{-i\omega_0 t}\left\{\frac{\gamma_0 - \gamma_{eff}}{\gamma_0 - 2\gamma_{eff}} e^{-(\gamma_0 - \gamma_{eff})t/2} - \frac{\gamma_{eff}}{\gamma_0 - 2\gamma_{eff}} e^{-\gamma_{eff} t/2}\right\} \qquad \text{(V.A.2.19)}$$

The first term corresponds to the coil excitation and dies out faster than the second term. This second term corresponds to the ion excitation. Note that the sign difference between the terms implies that at time t_0:

$$t_0 \approx \frac{2}{\gamma_0} \ln \frac{\gamma_0}{\gamma_{eff}} \qquad \text{(V.A.2.20)}$$

the current will go to zero because of the destructive interference between the two signals. (See Figure V.A.2.4a)

In the strongly coupled case, however, the two factored terms have the same damping ($\frac{\gamma_0}{2}$) and are split by $\delta \approx \sqrt{\gamma_0 \gamma_z}/2$. In that case, we get:

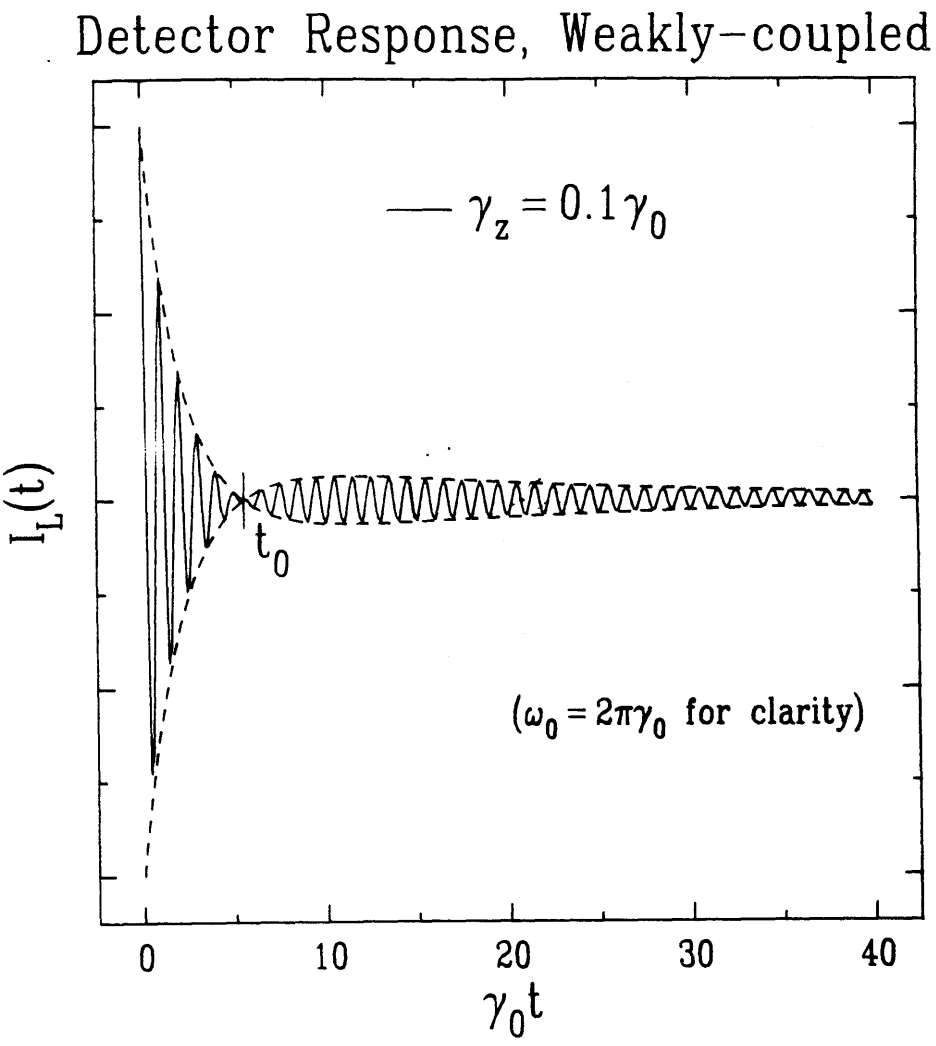

Figure V.A.2.4(a). The time-domain response of the detected current. At small times, the response is due primarily to the detector's excitation, while at later times, the response can be attributed to the decaying ions.

$$I(t) \sim e^{-i\omega_0 t} e^{-\gamma_0 t/4} \left[\cos\delta t + \frac{\gamma_0}{4\delta} \sin\delta t \right] \qquad \text{(V.A.2.21)}$$

Thus we see this current "sloshes" back and forth between the detector and the ions. The coil current goes to zero every $\frac{2\pi}{\delta}$ seconds, and damps at a rate slower than either the coil or the ions (if they were hooked up only to a resistor). (See Figure V.A.2.4b)

We also can use the admittance formalism to determine the motion of the trapped ions. As discussed in Section V.A.1, the velocity of the ions is proportional to the current passing through the *lc*-circuit in Figure V.A.3.2. Going back to the frequency-space representation, this current is:

$$I_{ion}(\omega) = [V_{drive}(\omega) - V_{det}(\omega)] \frac{i\omega c}{1 - \omega^2 lc} \qquad \text{(V.A.2.22)}$$

Plugging in Equation V.A.2.5 and grinding through the algebra yields:

$$I_{ion} = \frac{1 - \omega^2 LC + i\omega L/R}{(1 - \omega^2 L[C + C_{trap}] + i\omega L/R)(1 - \omega^2 lc) - \omega^2 Lc} i\omega c \; V_{drive} \qquad \text{(V.A.2.23)}$$

Going again to the resonance approximation gives a function similar to Equation V.A.2.6:

$$I_{ion} \approx \frac{i\omega_0^2 c}{2} \frac{\Delta_{LC}}{\Delta_z(\Delta_0 - i\gamma_0/2) - \gamma_0\gamma_z/4} V_{drive} \qquad \text{(V.A.2.24)}$$

except here the numerator is $\Delta_{LC} = \omega - \omega_{LC}$. The resonant structure is identical to that for the detector current—after all, these are the two oscillators which are coupled—but the ion motion has a zero at a substantially higher frequency, ω_{LC}. (I

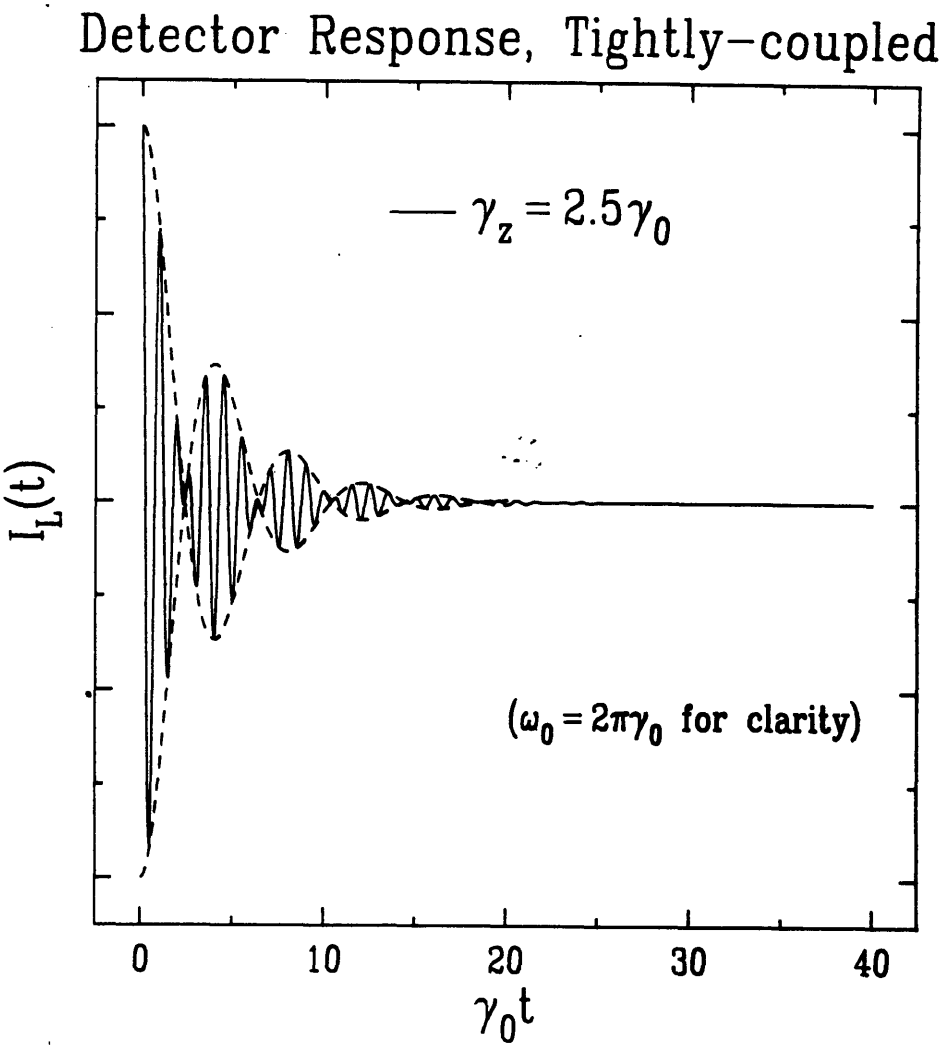

Figure V.A.2.4(b). The time-domain response of the detected current in the tightly-coupled case. The current can be seen to slosh back and forth between the ions and the detector.

have neglected the small imaginary part of the numerator because, near the detector and ion resonance, the real part dominates.) As argued above, this approximate zero arises because of the trap capacitance: at the resonance of the parallel LC-circuit, $V_{det} \approx V_{drive}$ and there is no voltage difference across the trap to drive the ions.

Another interesting effect of the capacitance is that the amplitude of the ions' response depends on the exact value of the trap capacitance. Since the ion resonance is far from the pure LC resonance, Δ_{LC} can be treated as a constant:

$$\Delta_{LC} \approx \omega - \omega_{LC} \approx \omega_0 - \omega_{LC} \approx -\frac{1}{2}\frac{C_{trap}}{C}\omega_0 \qquad \text{(V.A.2.25)}$$

Thus the ion response is directly proportional to the trap capacitance. Physically, this effect arises because right on the detector's resonance the voltage across the trap can be much larger than the drive voltage, boosted by the detector's peak. The degree of this enhancement depends, as in (V.A.2.25), on the trap capacitance.

In addition to the effect of C_{trap}, the response of each ion also can depend on the total number of ions in the trap. I will demonstrate this effect explicitly from equation (V.A.2.24) in the case that $\omega_z = \omega_0$. To determine the peak amplitude, we go to the time domain, factor (V.A.2.24) as we did before, and use the fact that the numerator is effectively a constant. The time domain behavior can be summarized in both coupling regimes by the relatively simple relation:

$$I_{ion}(t) \sim e^{-i\omega_0 t}\frac{1}{\sqrt{\gamma_0^2/4 - \gamma_0\gamma_z}}e^{-\gamma_0 t/4}\sinh\left[\sqrt{\gamma_0^2/4 - \gamma_0\gamma_z}\, t/2\right] \qquad \text{(V.A.2.26)}$$

In the strongly-coupled case, the argument of the radical will be negative, and we use the analytic continuation $\sinh iz = i\sin z$. In Figure V.A.2.5, I show calculated responses for several values of γ_z/γ_0. We can use (V.A.2.26) to determine the peak current and hence the peak ion velocity. The maximization is straight forward[1], though tedious, and the results are shown in Figure V.A.2.6. In that figure, I plot the peak ion velocity (normalized against one ion) versus the number of trapped ions. (For concreteness, I used the current setup, detecting N_2^+ ions.) Notice that this is an entirely different effect than the dependence of steady-state, driven response on the damping (as it would for a normal, decoupled oscillator.) Indeed, this effect should show up even for excitation pulses short compared to the damping time, for which the conventional wisdom holds that the damping can be neglected.

The final effect of the tuned circuit is that the peak ion response depends strongly on the frequency difference between the ion and coil resonance, $\omega_z - \omega_0$. I will demonstrate this effect explicitly when the ions are both far from the coil ($|\omega_z - \omega_0| \gg \frac{\gamma_0}{2}$) and have a small effective width. (Recall that the width decreases dramatically when the ions are moved from the coil, $\gamma_{eff} \approx \gamma_z \left[\frac{\gamma_0/2}{(\omega_z - \omega_0)}\right]^2$.) Under these conditions, we always will be in the weakly-coupled regime, and Equation

1. Interestingly, the peak in the ion velocity occurs at the same time as the detector current *minimum*.

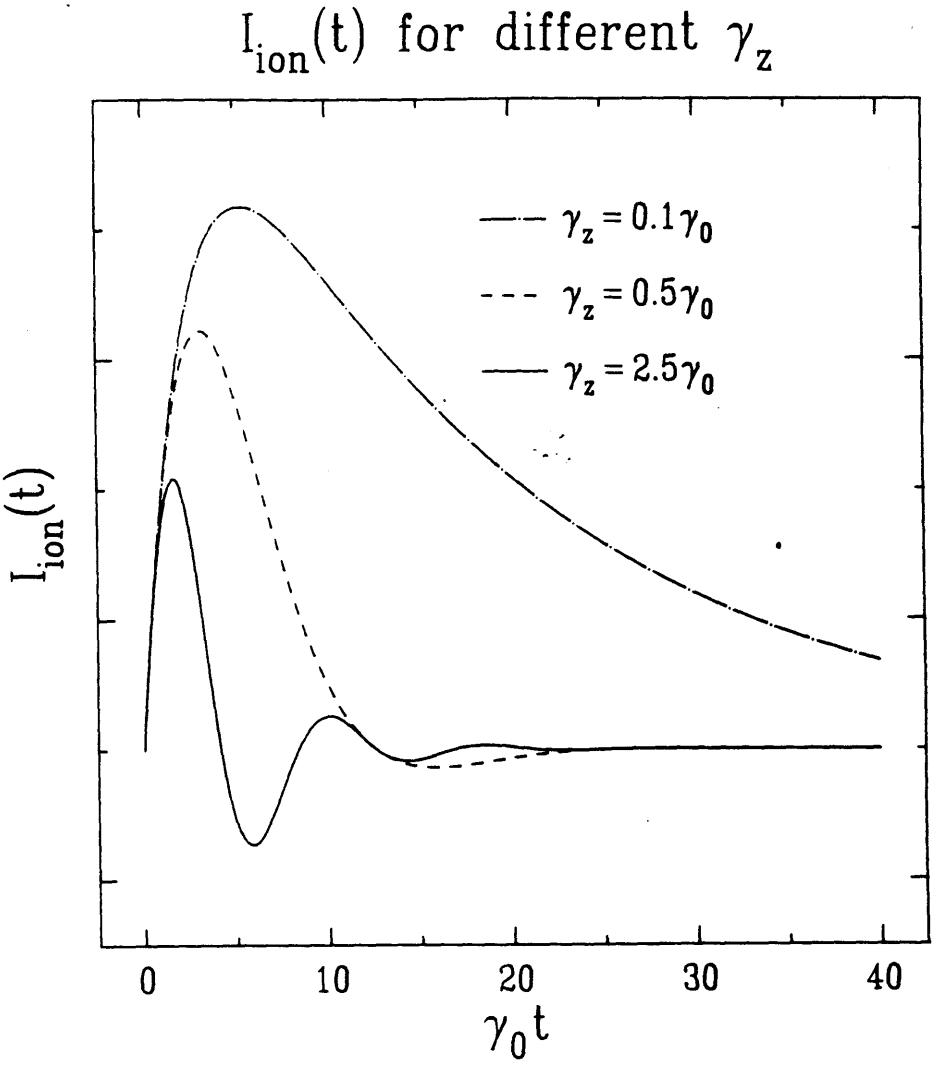

Figure V.A.2.5. The peak ion current for different ion dampings in the trap. This graph shows the current *perion*, demonstrating that the peak response decreases with increasing ion number.

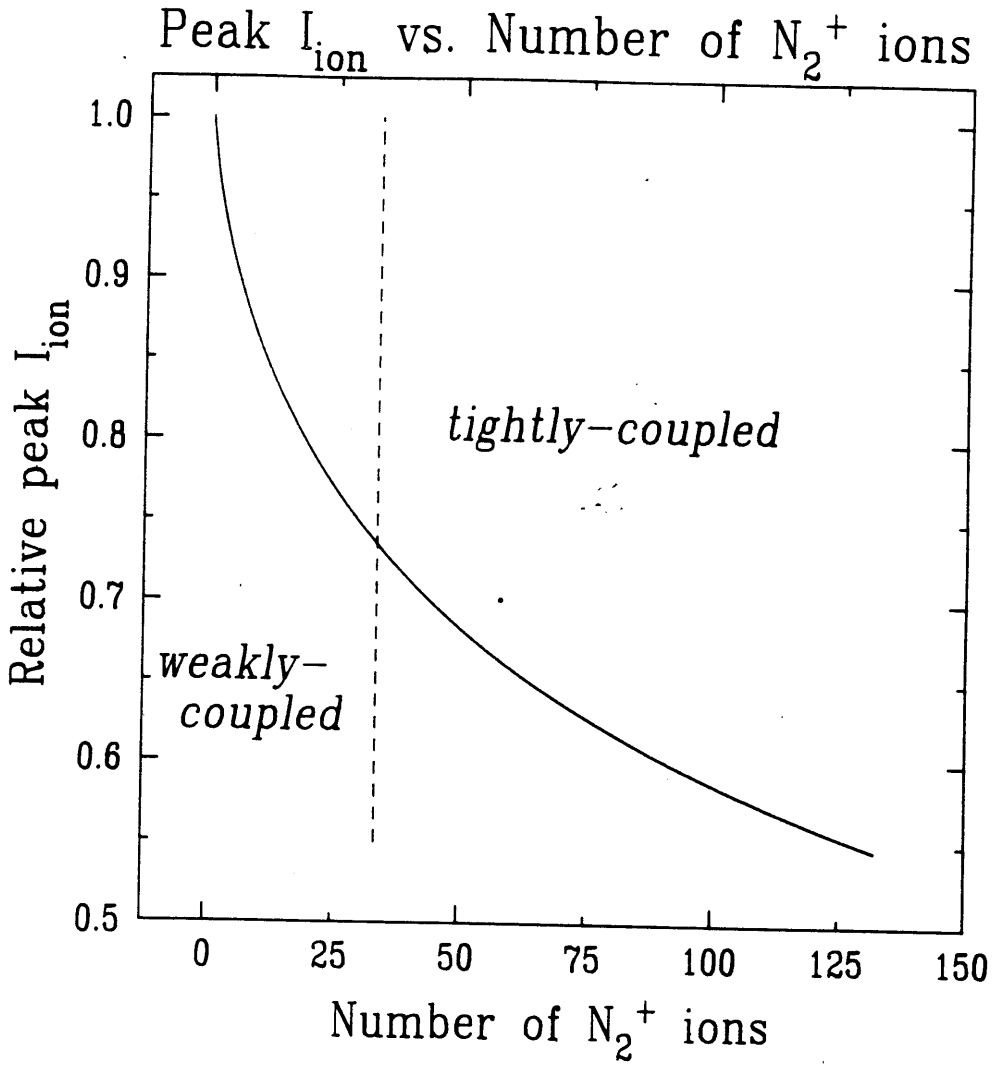

Figure V.A.2.6. The peaks from Figure V.A.2.5 for realistic conditions with our current detector ($\gamma_0 / 2\pi$ = 6.6 Hz and N_2^+ ions ($\gamma_z / 2\pi$ = 0.05 Hz per ion.

V.A.2.26 reduces to:

$$I_{ion}(t) \approx \frac{1}{\omega_z - \omega_0} e^{-i\omega_z t} e^{-\gamma_{eff} t/2} \qquad \text{(V.A.2.27)}$$

This expression is the current equivalent of the ion's motion, *not* the detected current. This frequency-dependent boost, as discussed before, arises from the large increase in the field that drives the ions when that field is produced near the detector's resonance.

To summarize this rather lengthy section, then, we can treat the trapped ions an a series *lc*-circuit, in parallel with the the trap capacitance. This treatment lets us make several predictions for the detector and ion responses to external drives. The capacitance of the trap and the presence of a *tuned* circuit (rather than a purely dissipative detector) profoundly affects the response of the ions to external drives. I have identified two regimes: in one the ions can be thought of as decoupled from the detector; in the other, the detector and ions are strongly coupled. In both regimes, the ions appear to short out the coil when they are resonant. The amplitude of the ion's response depends strongly on the trap capacitance and on the detuning of the ions from the coil. Additionally, the peak amplitude depends weakly on the number of ions in the trap, decreasing with larger clouds.

V.A.3 Response to CW and Pulsed Excitations

We can use the admittance formalism of the previous section to predict, quite simply, the frequency response from both a CW and a pulsed drive. A steady drive at a given frequency, ω_d, can be represented in frequency space as the sum of two

delta-functions at $\omega = \omega_d$ and $\omega = -\omega_d$ whose relative amplitude sets the phase. Therefore the response in the time domain will be at frequencies $\pm\omega_d$, and the physical amplitude of this response will be given by the magnitude of the admittance evaluated at $\omega = \omega_d$. (We ignore the overall phase of the result.) In our experiment, ω_0 is fixed (the resonant frequency of the detector cannot be adjusted except when the apparatus is warm), and we tend to keep ω_d fixed and vary ω_z by varying the trapping potential, Equation (II.A.2.5). Thus, we can compute the steady state response by evaluating equation (V.A.2.6) at ω_d. The resulting expression is a function of $\Delta_z = \omega_d - \omega_z$ with Δ_0 fixed at $\omega_d - \omega_0$. For large detunings, we find that:

$$I_L \rightarrow \frac{i\omega_0^2 C_{trap}}{2} \frac{1}{\omega_d - \omega_0 - i\gamma_0/2} \qquad \text{(V.A.3.1)}$$

and thus I_L is independent of the ions, as expected. As mentioned before, this current comes entirely from the capacitive feedthrough from the lower to the upper endcap of the trap, and we can use it to measure the Q of the detector as well as the trap capacitance.

However, on resonance ($\Delta_z = 0$) we find again that the detected current vanishes completely. We can compute the width of this dip, for example the full width at the half (squared-magnitude) maximum of the dip:

$$\Delta_z^{FWHM} = \gamma_z \left\{ \frac{\gamma_0^2/4}{(\omega_d - \omega_0)^2 + \gamma_0^2/4} \sqrt{1 + \frac{8(\omega_d - \omega_0)^2}{\gamma_0^2}} \right\} \qquad \text{(V.A.3.2)}$$

Thus, at fixed drive-coil detuning, the width of the dip is proportional to the

uncoupled ion damping regardless of the coupled widths computed in the previous section. In addition, when the drive is exactly resonant with the coil, $\omega_d = \omega_0$, the width of the dip yields exactly γ_z. Therefore, we can use the width of this dip to count the number of trapped ions. Figure V.A.3.1 shows the calculated coil response for several cloud sizes.

More frequently, however, we drive the ions with a short duration pulse, centered near ω_z, on the lower endcap. This pulse contains a fairly broad spectrum of frequencies, reasonably flat over a frequency band $\Delta\omega = 1/T$, where T is the duration of the pulse. If $\Delta\omega$ is much larger than the detunings and dampings in the problem, the resulting ion/detector spectrum will be identical to (V.A.2.6), multiplied by a constant proportional to the duration of the pulse, T. In other words, Equations V.A.2.6 and V.A.2.29 give the shape for the frequency spectra due to a pulsed drive and V.A.2.19,21 and 26 should give the time domain response.

For example, Figure V.A.3.2 shows the time-domain response of 25 N_2^+ ions excited by a single-drive pulse to about 8% of the trap size. As in the theoretical illustration in the previous section (weak-coupling, Figure V.A.2.4a), the faster decay in the early part of the transient can be attributed to the decay of the detector's excitation. The slower decay (after 0.1 s), on the other hand, came from the ions themselves. We can use the time at which the current envelope vanished ($t_0 = 0.11s$) to determine the damping of the ions, γ_z, as described in Equation V.A.2.20. This value, $\gamma_z / 2\pi = 1.2$ Hz, is consistent with more direct measurements of the ions' width.

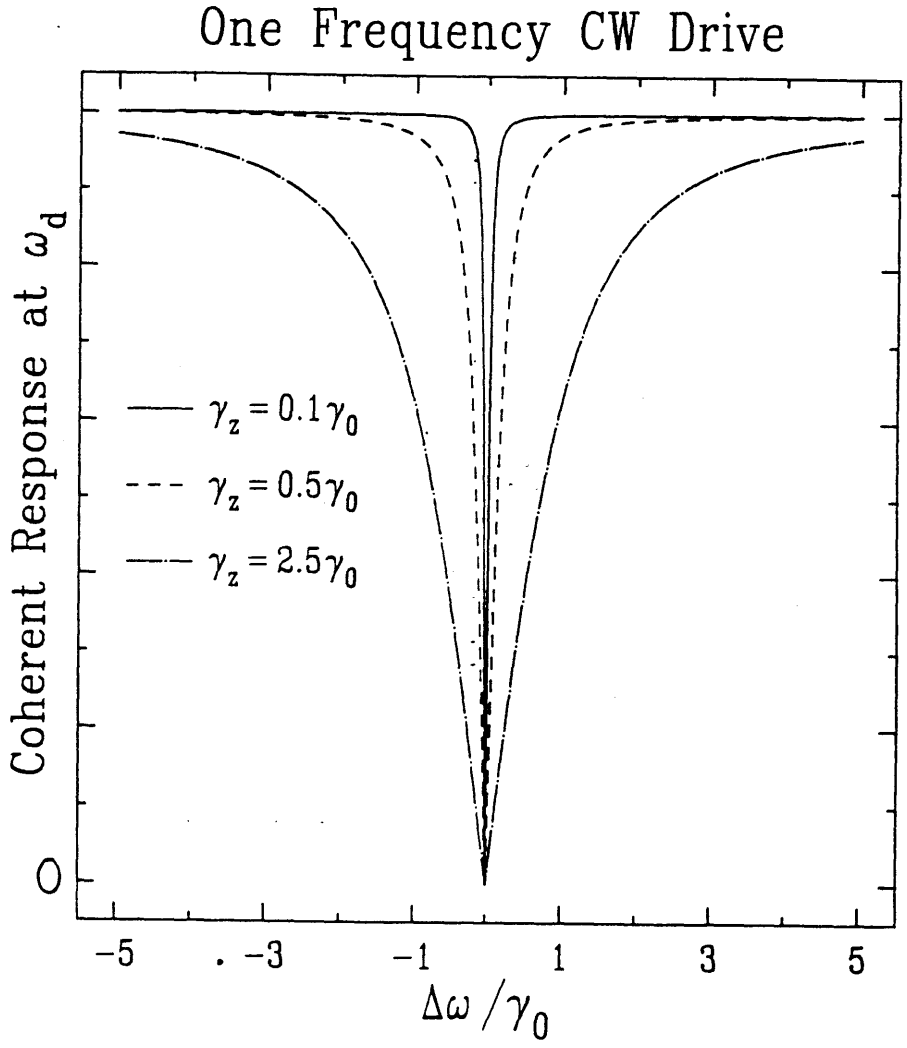

Figure V.A.3.1. The coherent response to a cw, one-drive excitation. Far from the ions' resonance, the response is constant, due to capacitive coupling from the lower to the upper endcap. On resonance, the detected signal vanishes as the resonant ions short-out the detector.

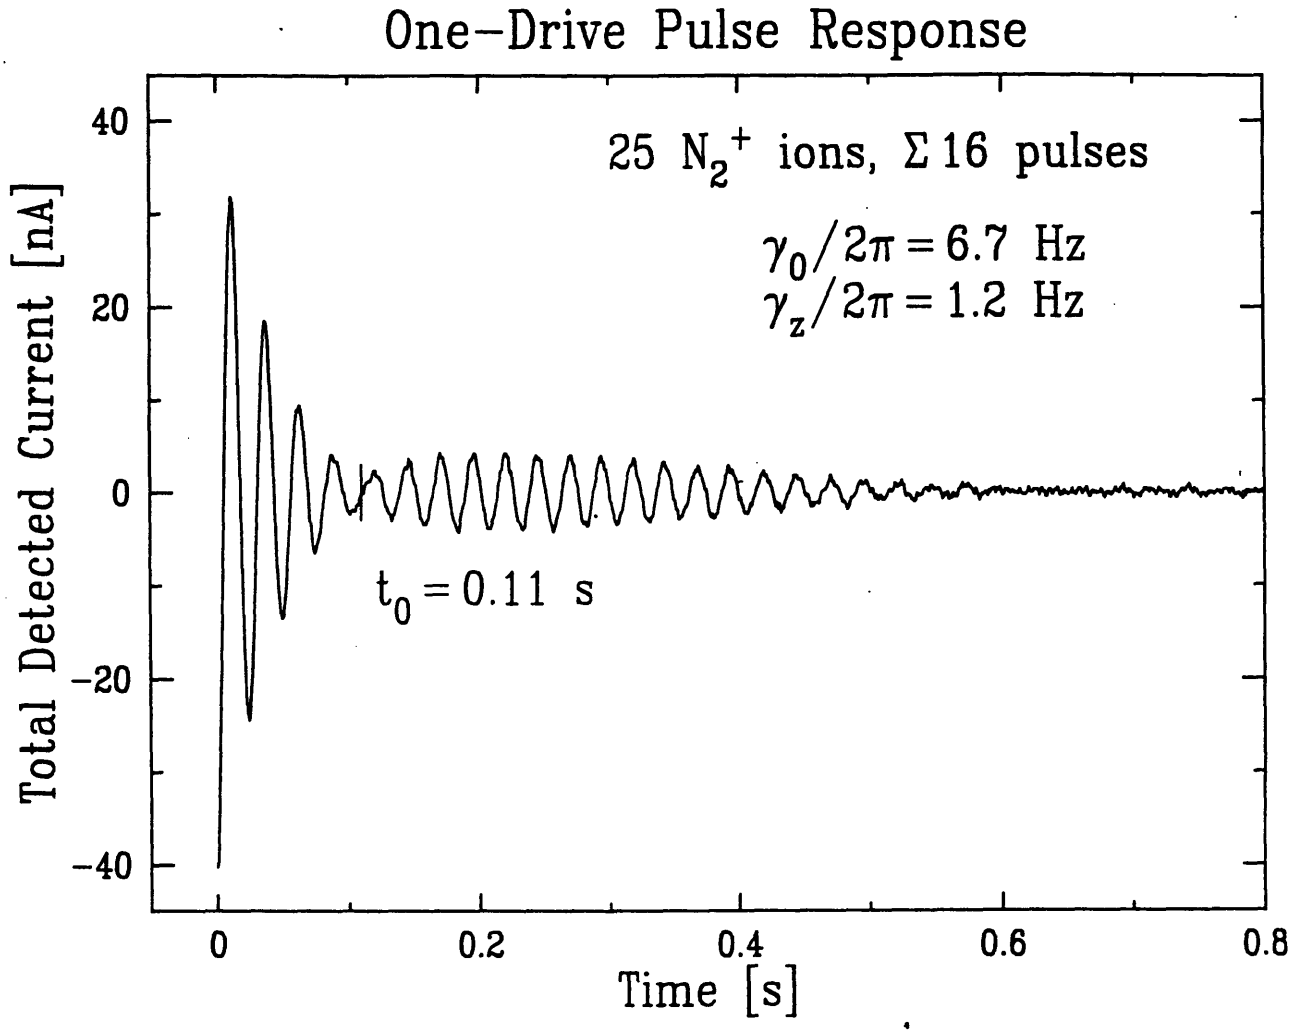

Figure V.A.3.2. Actual time-domain response for a small cloud of N_2^+ ions in the trap. These data are the sum of 16 transients.

In Figure V.A.3.2, the total current does not go completely to zero at t_0 because of residual anharmonicity. This anharmonicity can lead to amplitude—and thus time—dependent frequency shifts. This lack of reproducibility tends to blur some of the finer features when we add many scans together.[2]

As an example in the frequency domain, using the FFT, Figure V.A.3.3 shows typical ions and detector response to a one-drive pulsed excitation. The trap voltage was changed between each of the curves in order to move the 70 N_2^+ ions about 4 Hz. Each curve is the FFT of the sum of fifty identical, 4-second duration transient responses. Because of the different ion/detector detunings, the amount of the excitation should vary from about 10% of the trap size in the lowest curve to about 2.5% in the topmost, as described by Equation (V.A.2.27). However, in these curves we cannot see that effect directly. (However, see Figure V.B.2 in the next section.) The ion feature becomes noticeably narrower as the ions were moved further off the detector, a reflection of the decrease in the damping rate.

The effects described in the previous section—dependence of ion response on $\omega_0 - \omega_z$, C_{trap}, and number of trapped ions—can be directly observed by measuring the spectrum. As I will describe in the the next section, we have used the FFT on time-domain data to observe many these effects directly.

2. This figure is the sum of 16 independent scans.

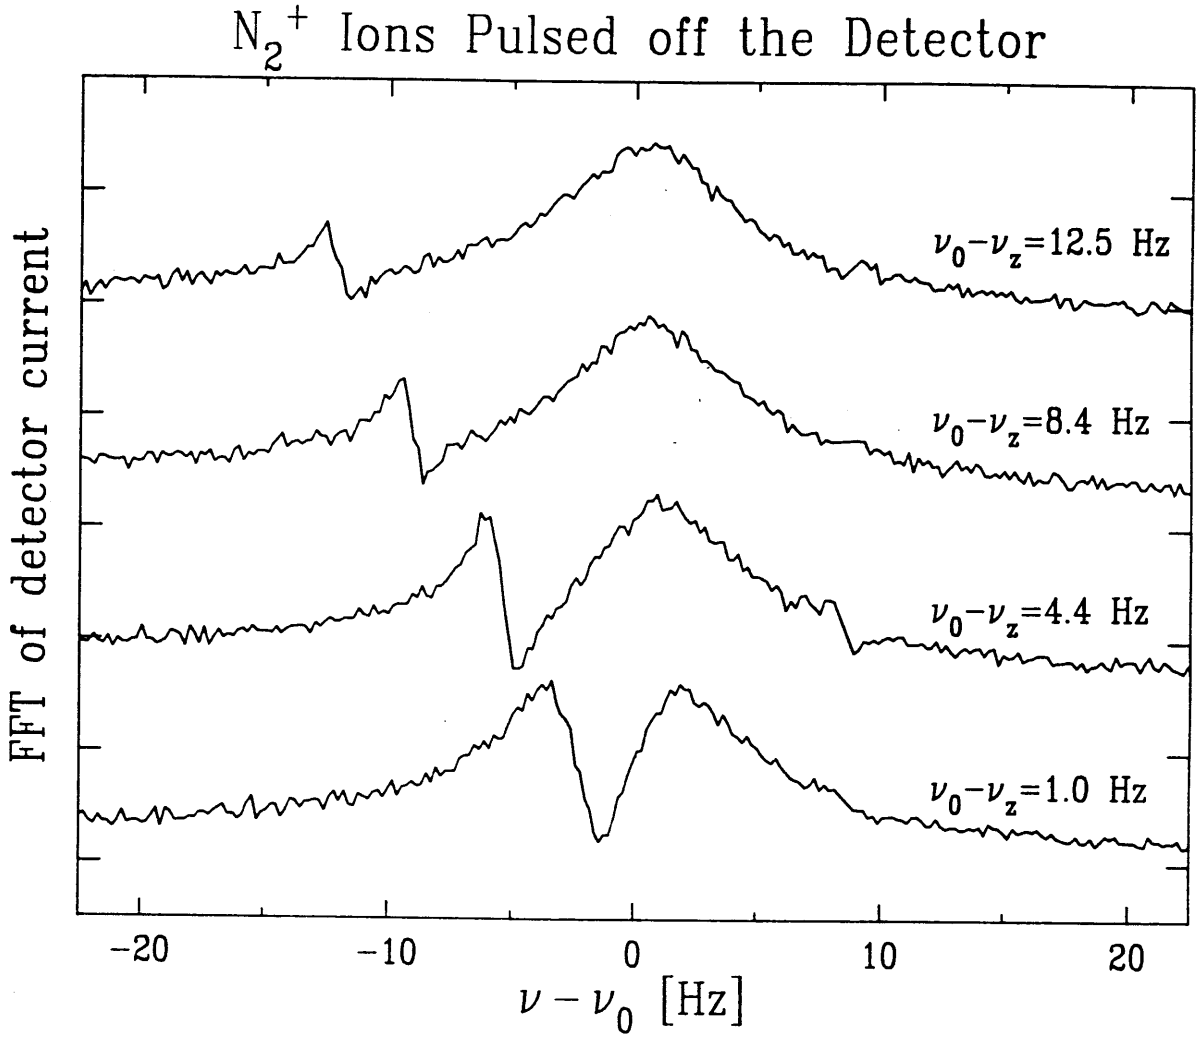

Figure V.A.3.3. Actual frequency-domain response for 70 N_2^+ ions. The ions are moved further from the detector as shown by the detuning labels on the right of the figure.

V.B One-Drive Pulse Results

In this section, we describe several experiments that used the one-drive, pulsed excitation scheme discussed in the last section. First, we will discuss a measurement in which we used the pulsed scheme to determine narrow ion decay widths by measuring their (rather large) resultant decay times. The second experiment shows the validity of some of the conclusions of the admittance method. In particular, we show that the ion excitation does, indeed, depend on the ion/detector detuning.

Both of these experiments used the coherent transient averager described in Chapter IV to improve the signal-to-noise. Although both also made extensive use of the FFT, the second measurement, in particular, shows some of the transform's subtleties. As we shall see, it must be used carefully to measure the amplitude of a decaying, sinusoid.

In the first experiment, we used the single-drive, pulsed excitations to measure the natural damping, γ_z, of one, two, and four ions, by measuring directly the decay of the ion signal as a function of time. In Figure V.B.1, we used two slightly different implementations of this concept to measure the damping time of several very few-ion clouds. In the upper two curves, we varied the length of time between the excitation and the beginning of our recording of the transient. When we fourier transform this response, the amplitude of its peak should depend on the delay time, T, like $e^{-\gamma_z T/2}$. For the bottom curve, we split up the time-domain transient response (16 seconds,

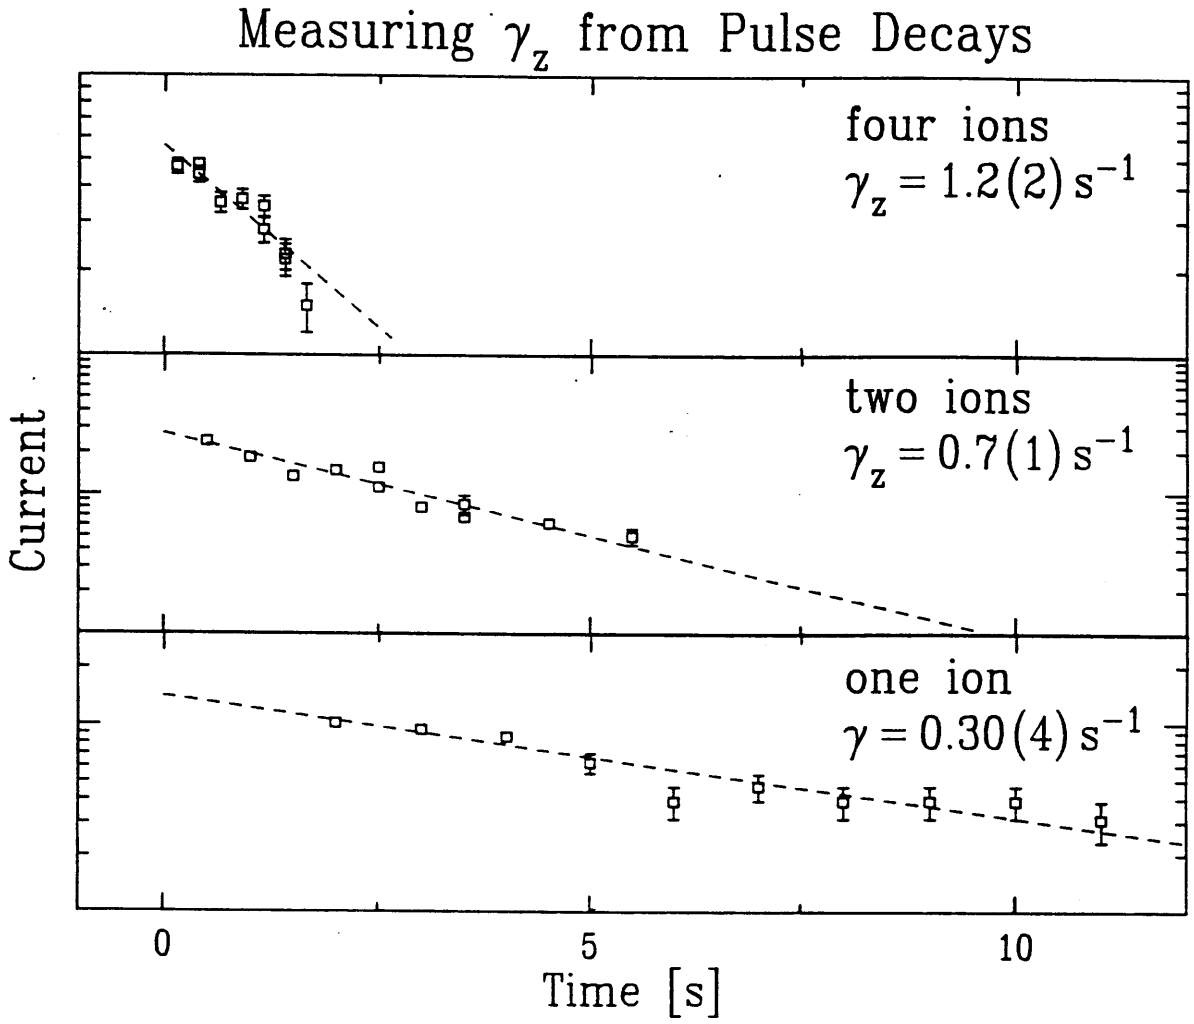

Figure V.B.1. Measuring γ_z by measuring the (exponential) time-dependence of the detected current.

total) into two-second segments and fourier transformed *each* of these segments. Again, the peaks should decay like $e^{-\gamma_z T/2}$, where T marks the beginning of each segment. In all three cases, we waited at least 0.2 seconds to allow the detector's excitation to die down completely. The single-ion decay gives $\gamma_z = 0.30(4)\ s^{-1}$, in excellent agreement with our a priori estimate, $\gamma_z = 0.33\ s^{-1}$. (To make this estimate, we inferred the impedance of the detector from its Q and inductance. The geometric factor, B_1, was estimated from [GAB83].) The two- and four-ion decays also agree well. (For the one- and two-ion curves, we knew the number of ions independently by a CW counting method, described below in Section D. The four ion curve, however, had no other, independent measurement of the number of ions.)

One of the more interesting predictions from the preceding section is that the pulsed response depends strongly on the detuning between the ions and the coil. (See Equation V.A.2.27.) To test this prediction, we excited a cloud of about 30 N_2^+ ions at many different ion/detector detunings and fourier transformed the resulting transient decays. The results, at various stages of analysis, are shown in Figure V.B.2. The raw data—FFT peak amplitudes—are given by the squares on the small, inset graph. The solid line on that graph is a prediction using the theory of FFTs from Chapter III and the theory of pulsed excitations from the last section. The dashed line on that graph shows the location and width of the detector (its amplitude is meaningless).

The predicted response comes from competition between three different effects. First, we expect that the ions' excitation will depend on their detuning from the

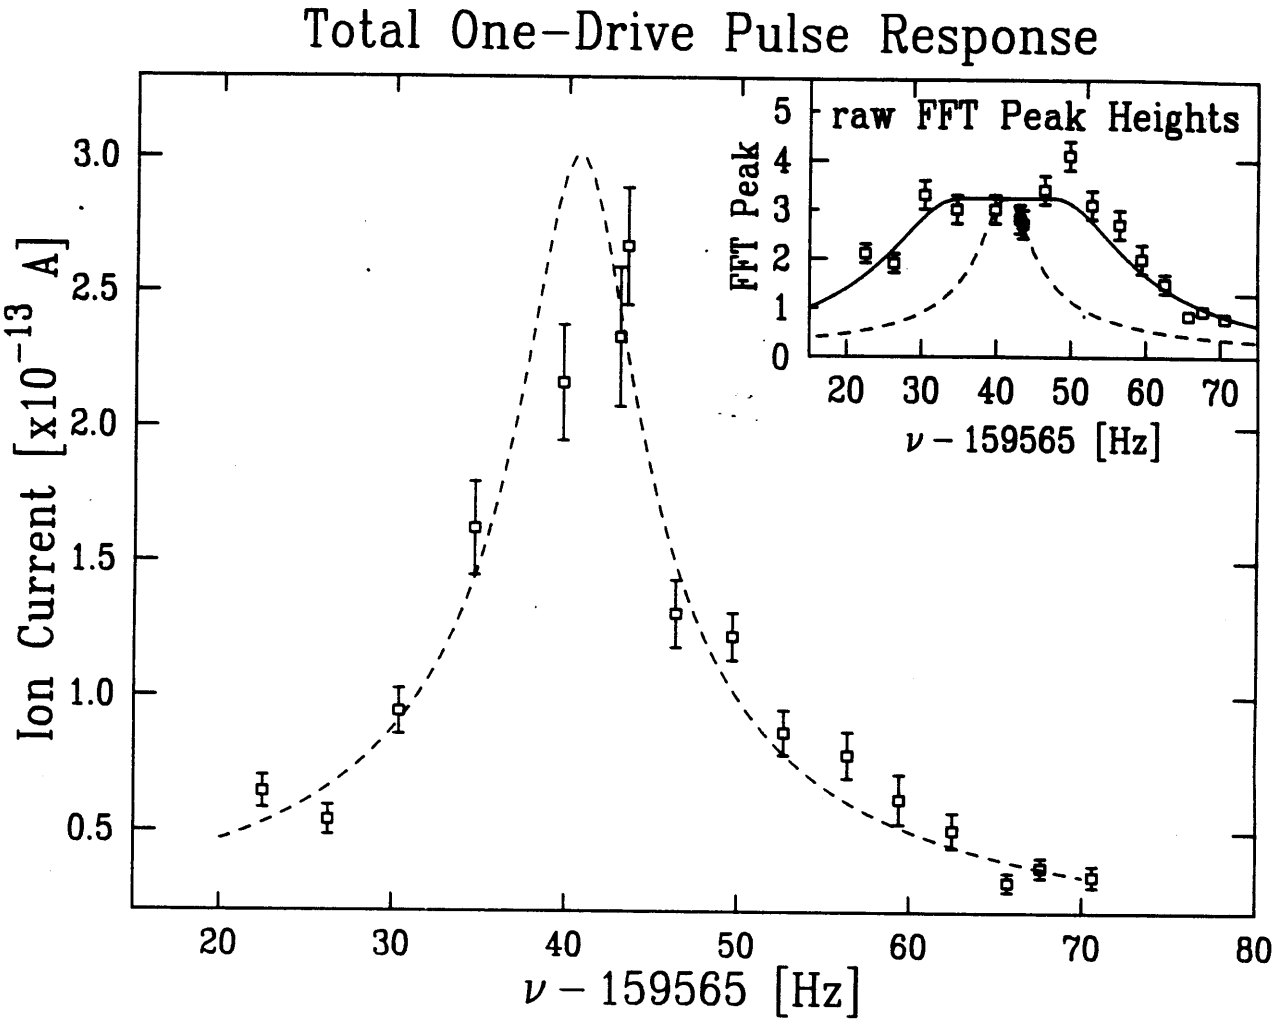

Figure V.B.2. Demonstration that the one-drive pulse response is strongly frequency-dependent. The larger graph shows (analyzed) amplitudes as a function of the ions' resonant frequency. The inset graph shows the raw, FFT data.

detector. Second, the gain of the detector decreases away from the peak of the coil. Third, the peak height of the FFT depends on the ions' damping. Thus, to demonstrate the *physical* effect (the first effect above), we must eliminate the two other *instrumental* effects.

Equation V.A.2.27 predicts that the ions' response to the peak should fall off like the magnitude of $\left(\omega_z - \omega_0 + i\gamma_0/2\right)^{-1}$. In addition, the gain of the detector falls of by an equal amount.[3] Therefore, we expect that the peak current induced by the ions should fall off like a lorentzian:

$$\frac{1}{(\omega_z - \omega_0)^2 + \gamma_0^2/4}$$

as we move the ions off the coil. Since the ion's damping shares this dependence (c.f. Equation (V.A.2.11) and the discussion below it.), we can say equivalently that the peak current should depend linearly on the ions' effective damping, γ_{eff}.

However, when we perform the FFT to go into the frequency domain, we alter this simple relationship. We can compute the peak FFT amplitude for a decaying sinusoid from the similar expression for a pure sine wave, Equation III.B.1.11. Inserting a complex frequency, $\omega + i\,\gamma/2$, into the Equations (III.B.1.5-11), we find that the peak amplitude in the FFT depends on:

3. As discussed above, this equality comes about because the detector itself causes the ions' frequency-dependence by boosting the electric field across the trap.

$$peak \sim \left\{1 - e^{\frac{-N\gamma}{2f_{samp}}}\right\}\frac{f_{samp}}{\gamma} \qquad \text{(V.B.1)}$$

where N is the number of points in the transform, f_{samp} is the sampling frequency, and γ is the damping. For very small damping, we find that the peak $\sim N/2$, in agreement with (III.B.1.6). However, for larger damping, $\gamma > 2f_{samp}/N$, Equation (V.B.1) shows that the peak height should depend inversely on the damping, γ.

In the experiment summarized in Figure V.B.2, $f_{samp} = 250$ Hz and $N = 1024$. From the argument above, then, peak heights for identical time-domain amplitudes will go like γ^{-1} as long as $\gamma > 0.5\,s^{-1}$. For 25 ions, this damping corresponds to about a 12 Hz detuning. However, in that region, as discussed above, the time-domain amplitude should behave like γ^{+1}. Therefore we expect the ions' FFT peak heights to remain constant: while the time-domain amplitude is *growing* like γ_z, the FFT peak is simultaneously *decreasing* by the same amount. For larger detunings, the ion signal decays negligibly during the time we observe the transient, N/f_{samp}. Thus the FFT no longer contributes a factor of γ, and, at these detuning, the FFT peaks quickly become smaller, decreasing like the detuning, squared.

The larger graph in Figure V.B.2, shows the peak current as determined from the frequency domain spectra. The FFT response therefore has been removed from the raw data. I have also removed the decrease in the signal due to the reduced gain of the tuned detector off its resonance. The resulting amplitudes clearly demonstrate the predicted behavior: the pulsed excitations fall off like the detuning of the ions from

the coil. The dashed line in the larger graph is a fit to this response. Only the overall amplitude was varied for this fit: the location and width of the resonance were held fixed.

In practice, however, we found the one-drive scheme most useful as a quick method for measuring, simultaneously, the ion and detector frequencies. For all but the smallest clouds,[4] we also got an immediate, though approximate, count of the number of ions from the width of the dip, as in Figure V.A.3.2.

However, it is difficult to use these techniques for precision measurements, for robust counting procedures, in the presence of anharmonicities, and so on. For example, the strong dependence of the ion/detector detuning makes it difficult to control the amplitude of the resulting excitation. In the presence of anharmonicities (which can shift the detuning *during* the excitation), the transient produced can be especially difficult to decipher. Therefore counting ions, for example, by looking for plateaus in amplitude of the decay transient proved insurmountably difficult. Looking for quantized steps in the ion width (as in Figure V.B.1?) also is a difficult proposition. The ions' damping depends very strongly on the ion/detector detuning (even more so than the resulting amplitude of the excitation), and thus, again, drifts and anharmonicity can make these measurements very difficult.

4. For $n < 10$, jitter and anharmonicity made the count a bit suspect.

In the next two sections, then, we will describe the theory and present results for a *two–drive* detection scheme. We will find that this two-drive scheme, though more complicated to implement, will provide far more robust techniques to count ions and measure trap systematics.

V.C Two-Drive Techniques

In this section, we describe techniques which overcome many of the difficulties of the one-drive schemes. Originally introduced [WED73] to eliminate the capacitive feedthrough from the drive to the detector, these techniques also prove more robust to disruptions by anharmonicity or contaminating ions and, in addition, make it easier to measure those disruptions. The "standard" two frequency drive scheme, which I will describe below, also provides additional ion counting methods and can be used to "lock" the ions to an external frequency source. With the proper choice of parameters, we also can use the two drives to produce many of the effects that a single drive can. However, the main advantage of using two drives is that the detector itself is not excited by the drives: it is only excited by the ions. This last feature obviously makes the interpretation of the data much simpler.

Two-Drive Green Function

The two-drive schemes use, in addition to the drive on the lower endcap mentioned in the last section, a second drive to modulate the trapping potential. That is, the resonant frequency becomes time-dependent:

$$\omega_z^2 \rightarrow \omega_z^2(1 + \lambda\cos\omega_{LO}t) \qquad \text{(V.C.1)}$$

where ω_{LO} is the frequency of the modulation and λ represents its depth. (The subscript "LO" stands for "local oscillator," a term borrowed from RF-mixer terminology.) The simple harmonic oscillation equation of motion, (III.A.1.1), thus becomes a driven mathieu equation [MCL47]:

$$\ddot{z} + \gamma_z \dot{z} + \omega_z^2(1+\lambda\cos\omega_{LO}t)z = -\frac{e B_1}{2mz_0}V_{drive} \qquad \text{(V.C.2)}$$

There are two simple ways we can make ω_z periodic. On the one hand, if we add a small oscillating drive, V_{LO}, to the static trapping potential on the ring, V_{DC}, then the resonant frequency term in the equation of motion becomes:

$$\omega^2 = -\frac{eV_{ring}}{md^2} = \frac{eV_{DC}}{md^2}(1 + \frac{V_{LO}}{V_{DC}}\cos\omega_{LO}t) \qquad \text{(V.C.3)}$$

and thus, referring to Equation V.C.2 above,

$$\lambda = \frac{V_{LO}}{V_{DC}} \qquad \text{(V.C.4)}$$

On the other hand, we could put the additional drive on the lower endcap. As discussed in Section II.B, when there are three different potentials on the ring and endcaps, the part of the potential that is symmetric under z $\rightarrow$ -z will trap the ions, while the anti-symmetric part will shift the center of the ion cloud. Thus, when an additional drive, $V_{LO}\cos\omega_{LO}t$, is put on just the lower endcap, half of it will drive the ions, and half will shift the resonant frequency:

$$\omega^2 = \frac{eV_{DC}}{md^2}(1 - \frac{1}{2}\frac{V_{LO}}{V_{DC}}\cos\omega_{LO}t) \qquad \text{(V.C.5)}$$

and thus,

$$\lambda = -\frac{1}{2}\frac{V_{LO}}{V_{DC}} \qquad \text{(V.C.6)}$$

Thus, the additional drive on the lower endcap is similar to a drive on the ring, though only half as effective. (The change in sign makes little difference.) The half of the drive that shifts the cloud is so far from resonance that we can neglect it.

(In passing, note that the one-drive schemes from the previous section also modulate the trapping potential. The damping, γ_z is modulated, too, because it, arises from a voltage on only one endcap. However, the amplitudes of these voltages are much smaller than the ones we will be discussing below and, in addition, are at frequencies at which their effects will be small.)

The additional term, $(\lambda\omega_z^2\cos\omega_{LO}t)z$ can have dramatic consequences for the ions. Because this term destroys the time-origin (shift) invariance of the system, it creates several new effects. Most importantly, the destruction of shift-invariance allows the generation of harmonics and subharmonics, and introduces possible regions of unstable growth in the ions' response.

The simplest case of instability occurs at $\omega_{LO} = \omega_z/2$. Neglect damping and assume that there is some small amount of ion motion at ω_{LO}, $z = A\cos\omega_{LO}t$. Then the term $(\cos\omega_{LO}t)z$ becomes $A\cos^2\omega_{LO}t = A(1 + \cos\omega_z t)/2$. This additional term

thus looks like a resonant acceleration at ω_z and, since it is proportional to A, it can grow exponentially. (Arnold [ARN78] gives a good treatment of these unstable islands, including damping.) Although such "resonances" are used in some Penning trap experiments, we have avoided this method since the resulting amplitude is usually set by trap or even detector anharmonicities, and thus not wholly time independent.

The drive at ω_{LO} also can produce sidebands. For example, ion motion at an arbitrary frequency, $z = A\cos\omega t$, yields, for the additional term, $A\dfrac{\cos(\omega+\omega_{LO})t + \cos(\omega-\omega_{LO})t}{2}$ which to the ions acts like additional drives at $\omega \pm \omega_{LO}$. These forces induces motion at those frequencies, which, in turn, induces additional motion at $\omega \pm 2\omega_{LO}$, and so on.

Using this simple approximation technique, we can estimate the response due to two drives, $V_d\cos\omega_d t$ and $V_{LO}\cos\omega_{LO} t$ when $\omega_d + \omega_{LO} \approx \omega_z$; that is, when the two drives sum to the ion's resonance frequency. We also use the simplifying assumption that $\omega_{LO} \ll \omega_d, \omega_z$, and, in addition, $\omega_{LO} \gg \gamma_z$. Under these assumptions, ω_d will be quite far from resonance, and thus we neglect damping and write down the simple harmonic oscillator response of the ion due to the non-resonant drive at ω_d:

$$z = -\frac{e B_1}{2mz_0}\frac{1}{\omega_z^2 - \omega_d^2}V_d\cos\omega_d t \qquad \text{(V.C.7)}$$

We now concentrate on the part of $(\omega_z^2\lambda\cos\omega_{LO} t)z$ oscillating near ω_z since that is the part that is "resonant" and thus will have a large effect:

$$(\omega_z^2 \lambda \cos\omega_{LO} t)\, z \approx \frac{1}{2}\omega_z^2 \lambda \cos(\omega_{LO} + \omega_d)t \cdot \frac{eB_1}{2mz_0} \frac{1}{\omega_z^2 - \omega_d^2} V_d \qquad \text{(V.C.8)}$$

This term looks like an acceleration at $\omega_{LO} + \omega_d$ caused by a voltage:

$$V_d' = \frac{1}{4}\lambda \frac{\omega_z}{\omega_{LO}} V_d \qquad \text{(V.C.9)}$$

where I have used $\omega_z^2 - \omega_d^2 \approx 2\omega_z \omega_{LO}$. Thus, when this approximation holds, these two drives act just like a single drive with its amplitude reduced by $\beta/2$, where $\beta = \frac{1}{2}\lambda \frac{\omega_z}{\omega_{LO}}$, the traditional "modulation index" of FM, the ratio of the maximum frequency deviation to the modulating frequency. [HOH80]

Although we could, in principle, continue this approach to higher orders, it becomes increasingly awkward. As as alternative, we can generalize the Green function of the simple harmonic oscillator to include the effects of ω_{LO}. Since this approach was used by Brown and Gabrielse [BRG86], I will only sketch their results.

They make a generalization of the standard harmonic oscillator Green function:

$$G_z(t-t') = \frac{\theta(t-t')}{\omega_z} e^{-\gamma_z (t-t')/2} \sin \omega_z (t-t') \qquad \text{(V.C.10)}$$

by replacing ω_z with $\overline{\omega}_z \equiv \omega_z (1 + \lambda \cos\omega_{LO} t)^{1/2}$:

$$\overline{G}_z(t,t') = \frac{\theta(t-t')}{(\overline{\omega}_z(t)\overline{\omega}_z(t'))^{1/2}} e^{-\gamma_z (t-t')/2} \sin \int_{t'}^{t} d\bar{t}\, \overline{\omega}_z(\bar{t}) \qquad \text{(V.C.11)}$$

This expression is their equation 3.52. It is correct for small $\left[\frac{\lambda\omega_{LO}}{\omega_z}\right]^2$, which is almost always the case for us, even when $\beta \gg 1$. Notice that $\overline{G}_z$ does not depend

solely on $(t-t')$, but depends, explicitly, on the value of t. This fact makes fourier analysis much more difficult and thus, for the moment, we remain in the time domain.

The key to their solution lies in expressing the sin integral as a double sum:

$$\sin\int_{t'}^{t} d\bar{t}\,\bar{\omega}_z(\bar{t}) \approx \sum_{n,n'} J_n(\beta)J_{n'}(\beta)\sin\left\{(\omega_z'+n\omega_{LO})t - (\omega_z'+n'\omega_{LO})t'\right\} \quad \text{(V.C.12)}$$

where

$$\omega_z' \equiv \omega_z\left(1-\frac{\lambda^2}{16}\right) \quad \text{(V.C.13)}$$

Thus the second drive has two important effects. First, it causes a small, but measurable, shift in the natural resonance frequency, ω_z. The second effect arises from all those sin terms. In a Green function, a term $\sin(\omega_A t - \omega_B t')$ means that a drive at ω_B will induce a steady-state response at ω_A. Thus, the double sum expresses the sideband structure ω_{LO} makes possible.

In our experiment, when we drive the ring at ω_{LO}, we often drive the endcap at $\omega_d \approx \omega_z - \omega_{LO}$ and detect near ω_z. Thus we are interested in the term of (V.C.12) that looks like $\sin(\omega_z' t - (\omega_z' - \omega_{LO})t')$; that is, the n=0, n'=1 term. In particular, then, we have an approximate steady-state response:

$$z(t) = J_0(\beta)J_1(\beta)\left\{\mathrm{Re}\left(\frac{eB_1}{4mz_0\omega_z}\frac{V_d}{(\omega_z-(\omega_d+\omega_{LO})+i\gamma_z/2)}e^{-i(\omega_d+\omega_{LO})t}\right\}\right. \quad \text{(V.C.14)}$$

The term inside the { } is the "normal" response of the ions to a drive $V_d\,e^{-i(\omega_d+\omega_{LO})t}$

Thus the total effect of the second drive is to reduce V_d by the amount $J_0(\beta)J_1(\beta)$.

For small β, $J_0(\beta) \approx 1$, $J_1(\beta) \approx \beta/2$, and thus we recover the earlier approximation, Equation V.C.9. However, for stronger drives we find that the amplitude of the response begins to decrease with increasing V_{LO}, and eventually even goes to zero at $\beta \approx 2.4$, the first zero of J_0. (See Figure V.C.1) We have used both this effect and the much smaller shift in ω_z' with λ check the amplitude calibration of the ring drives.

For example, Figure V.C.2 shows a series of measurements in which we used the Bessel function dependence in the ring drive to verify our voltage calibrations. We loaded a rather large cloud (about 400 N_2^+ ions) into the trap. After we used the one-drive, pulsed technique (described in the previous sections) to determine their natural resonance frequency, we moved the ions onto the detector, $\omega_z = \omega_0$. We then set up the two-drives to excite the ions about 100 Hz off their resonance; that is, we set the lower endcap drive at $\omega_d = \omega_z - \omega_{LO} - 2\pi(100\text{ Hz})$. We detuned this far from resonance to simplify the interpretation of the results: for our apparatus, the ring drive required to reach the first zero of J_0 ($\beta \approx 2.4$) should shift the ions' resonance (Equation V.C.13) by about 9 Hz. By driving the ions 100 Hz off resonance, this shift caused less than a 10% change in the ions' response. Using this method, the calibration agreed to about 15% with our room-temperature measurements, Section IV.F.

With fixed local oscillator drive, the ions' response to the drive at ω_d should look just like a harmonic oscillator excitation. In Figure V.C.3, we show a typical

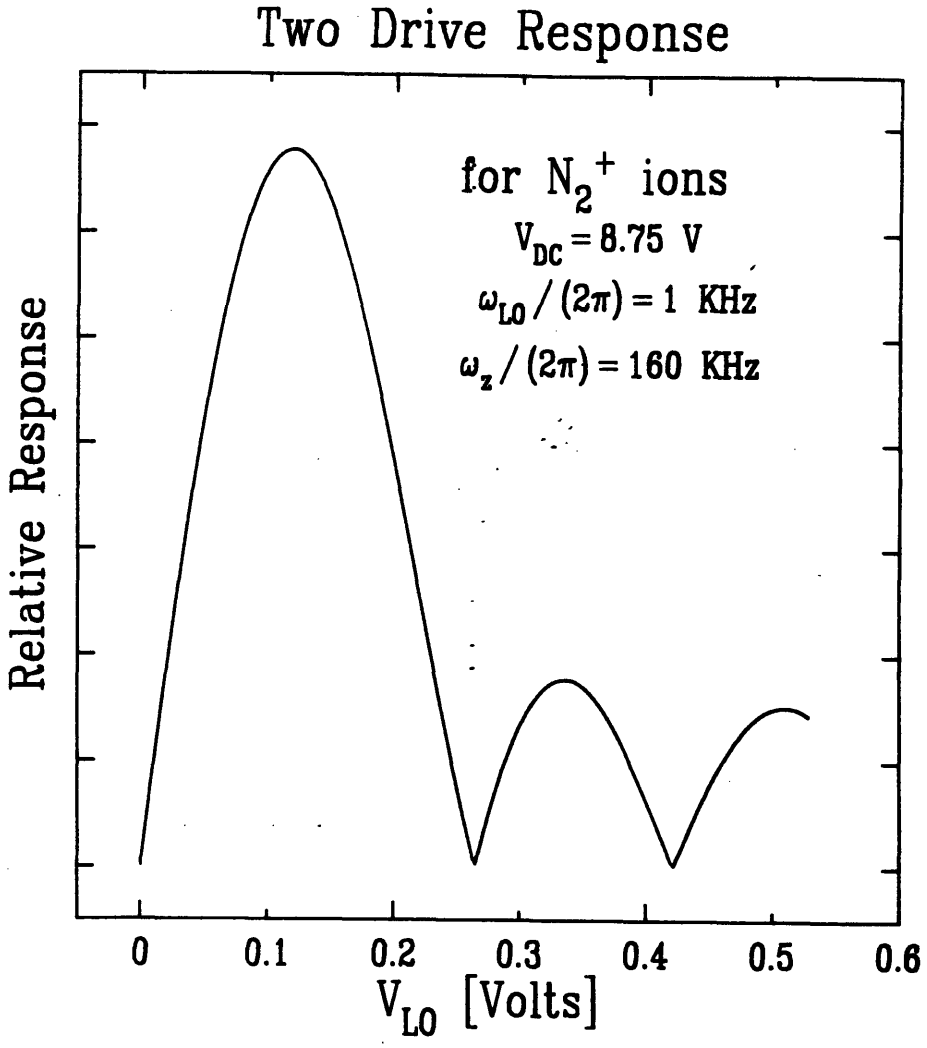

Figure V.C.1. Bessel-function dependence for the resulting ion amplitude when two-drives are used to excite the ions. (Theory)

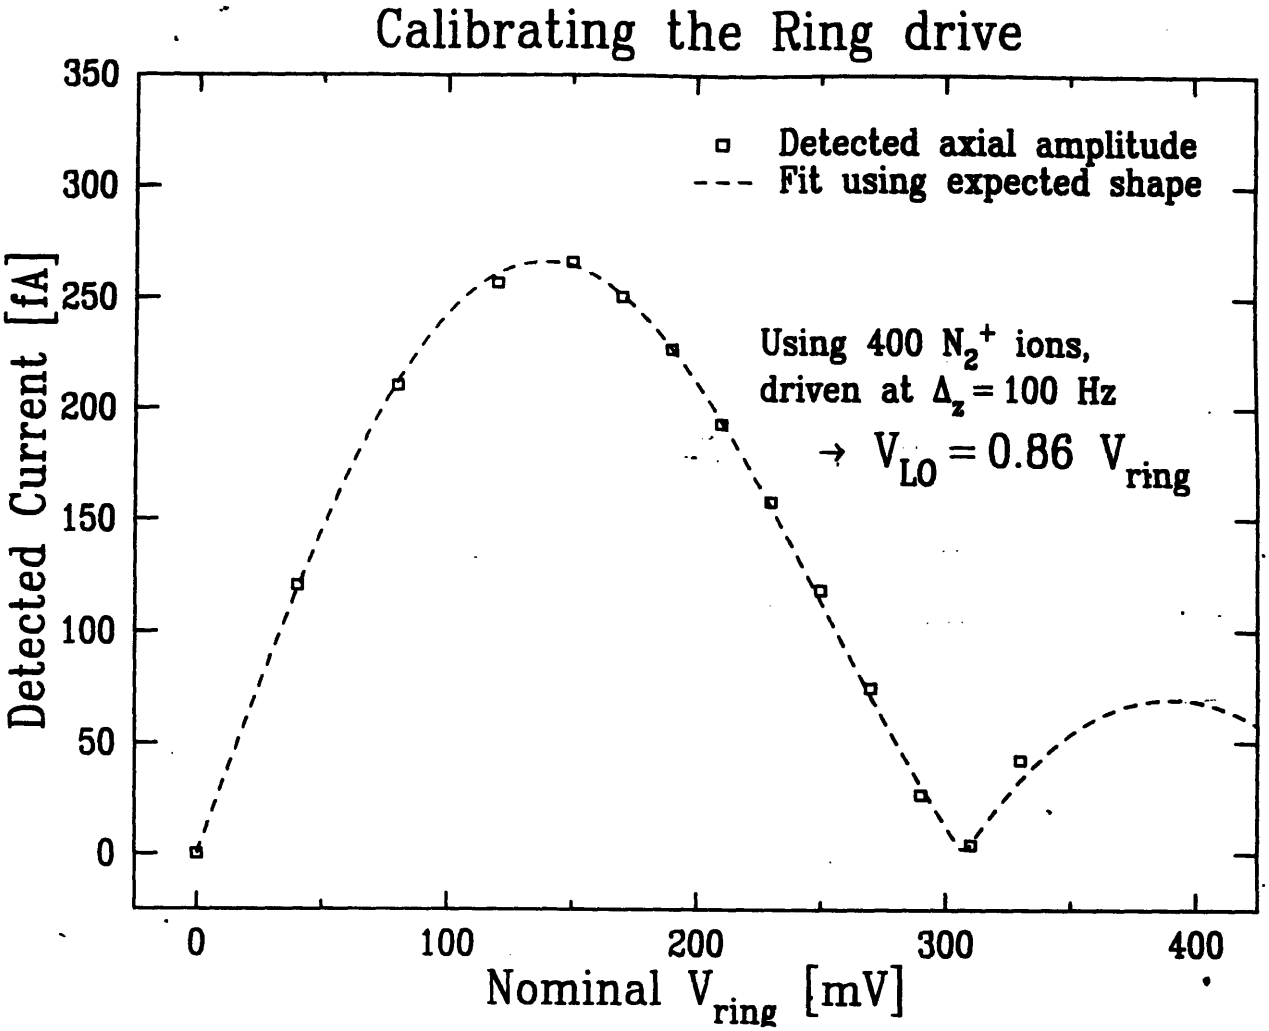

Figure V.C.2. Using the dependence in Figure V.C.1 to calibrate the ring drives. In this case, there was a 15% difference between our room-temperature measurements and the cryogenic calibration.

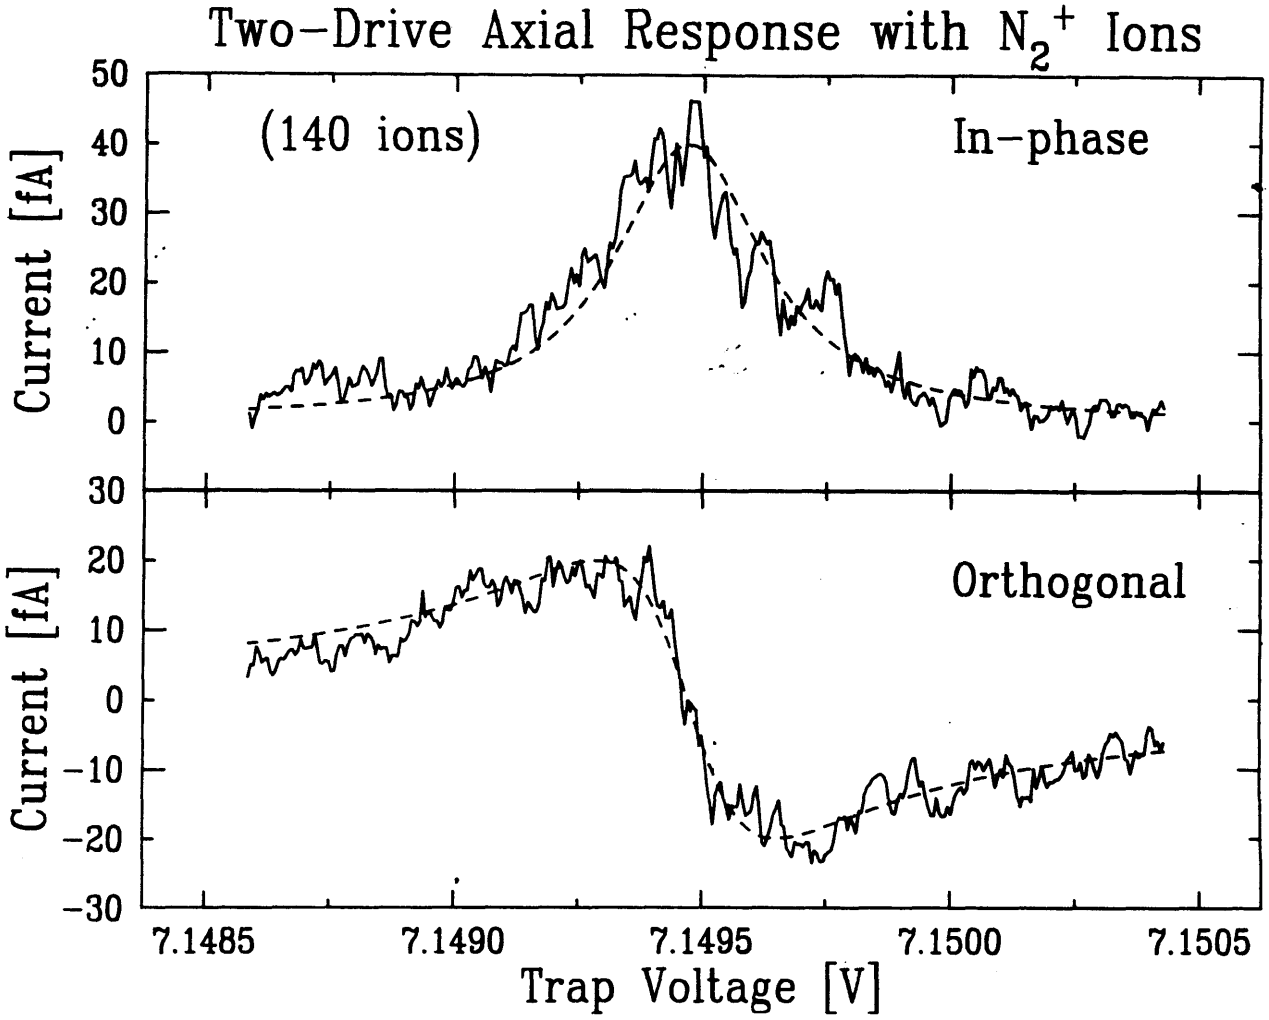

Figure V.C.3. Axial response for 140 N_2^+ ions, demonstrating normal, harmonic oscillator response.

example of a swept, CW resonance. Using the phase-sensitive detection technique described in Section IV.E.3, we measured both the real ("In-phase") and imaginary ("Orthogonal") parts of the response, (V.C.14). From the width of this resonance, we conclude that we had about 140 ions in the trap. (These data were taken with our older, Q=12000 detector. For that detector, roughly 40 ions cause one Hz of line-width.) We drove the ions very lightly for this sweep; their maximum excursion was about 30 μ at the peak of the resonance.

V.D Two-Drive, CW Results

The two-drive scheme avoids many of the complications which arose from a single drive. For example, since neither drive is sufficiently close to the detector resonance, the ion does not see a drive significantly boosted by the detector response. In addition, when we drive at constant frequency and sweep ω_z, (similar to our discussion in Section V.A.3 for the one-drive case) the response will be identical to an oscillator with resonance frequency ω_z and damping γ_z, even though the *true* damping is still γ_{eff} given in Section V.A.2. Thus the two frequency CW scheme lets us ignore, almost entirely, the subtle issues brought up in the last sections. With two drives, we can treat a CW resonance as if it were caused by a drive $J_0(\beta)J_1(\beta)V_d$ at a frequency $\omega_d + \omega_{LO}$, measured by an entirely uncoupled detector.

Ion Counting

The two-drive scheme provides two additional methods for counting the number

of ions in the trap. The first method, sometimes called the "Milliken method," is technically impossible using one-drive schemes. This method provides the most convincing, and, for that matter, the most robust ion-counting technique. As we see from (V.C.14), far from the ions' resonance, z(t) becomes independent of γ_z, and thus each ion acquires a velocity independent of the total number of ions. Hence, the total induced current will be proportional to the total number of ions. By selectively expelling ions from the trap, we therefore can observe quantized steps in the output current. Like Milliken, then, we can count the ions; here, without exact knowledge of the trap parameters, like z_0 and B_1, or the drive and detector gains. This counting method is thus far less model-dependent than any of the counting schemes that depend on measuring an *absolute* value of γ_z.

In Figure V.D.1, we show the first series of single-ion steps we observed with our apparatus.[5] We first loaded a small number of N_2^+ ions into the trap. Each plateau in the figure was made with the same excitation drives. (We increased the drive V_d from zero in about 30 seconds to avoid transient effects. We reversed the ramp at the end of each plateau for the same reason.) We drove the ion(s) 15 Hz off their resonance, to about 20% of the trap size at the peak. Between each of the plateaus, we lowered the potential on the lower endcap in order to force the ions perilously

5. March 2, 1988

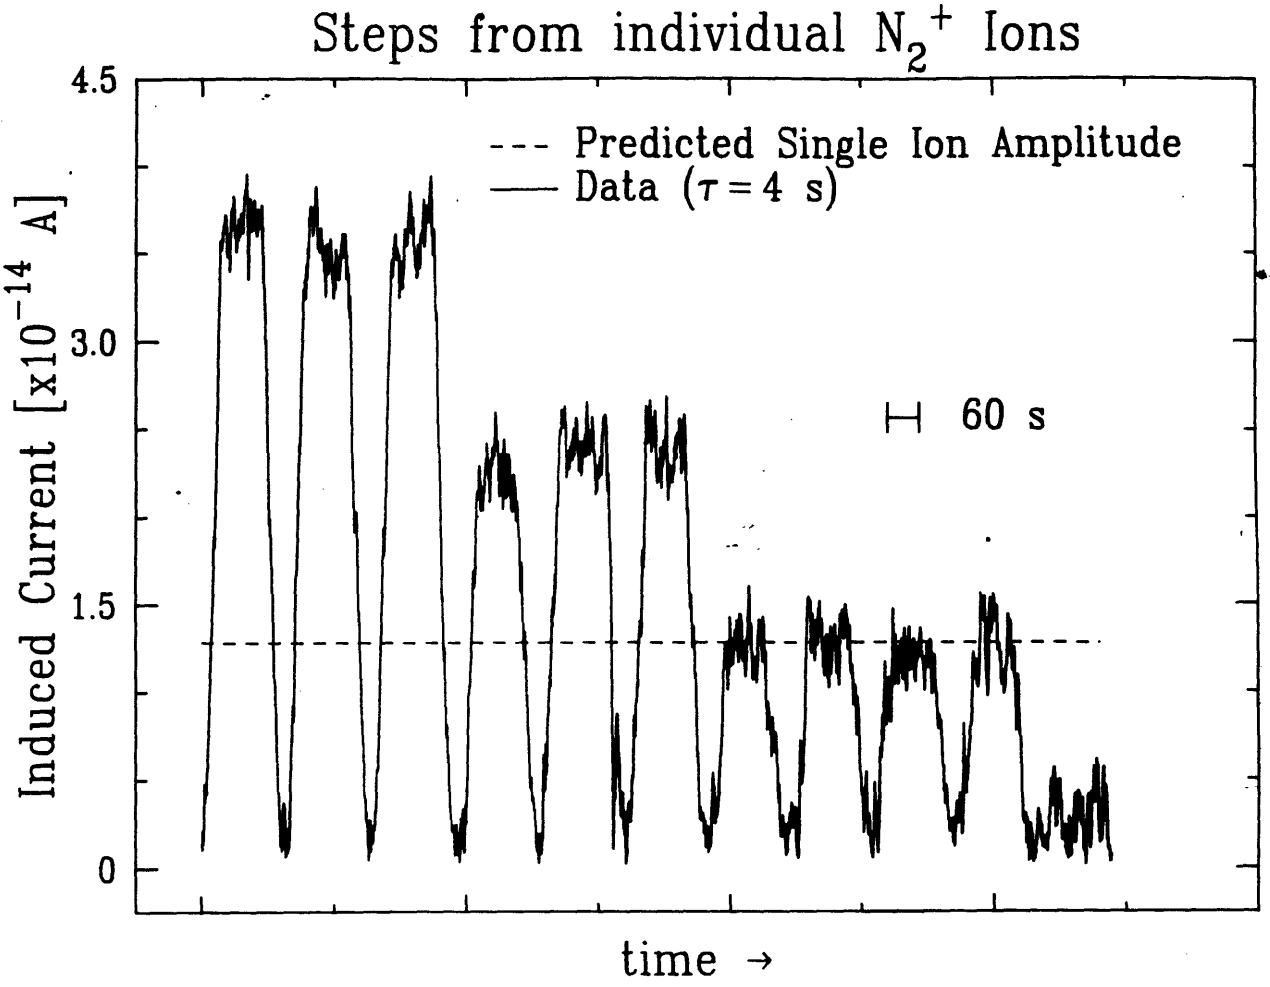

Figure V.D.1. Steps in the axial signal as one ion after another is expelled from the trap. The ions were driven to 20% of the trap size.

close to the lower endcap, (See Section II.B) and then "cooled" the remaining ions by resting on the detector for roughly 60 seconds. We then repeated the ramp-wait-ramp procedure. By slowly decreasing the endcap voltage between plateaus, we could expel one ion at a time. In Figure V.D1, we lost one ion between plateau 3 and 4, another between 6 and 7, and the final ion after plateau 10.

As shown on the figure, our estimates for the trap size, drive amplitudes, detector gain, etc., were fairly good, and the size of the resulting, single-ion signal agreed within its error to our a priori prediction. To our knowledge, these data represent the largest m/e yet detected as a single ion in a Penning trap, and brings to two the number of research groups that have detected a single, trapped ion electronically.[6]

The second counting method, the analogue of the CW, one-drive method, uses the width of a swept resonance. This resonance has a peak whose width is γ_z, and thus is proportional to the number of ions in the trap, as in Equation V.A.1.7. We used this method, for example, to determine the number of ions in the cloud from Figure V.C.3.

To convince ourselves that the widths we measured indeed were due to detector-induced damping, we performed some measurements by moving the sum of

6. Several other groups have performed single-ion experiments using lasers.

the drives off the detector resonance. Under these conditions, the swept response should become narrower because γ_{eff} becomes smaller off resonance (Equation V.A.2.11).

In Figure V.D.2, we have trapped about 140 N_2^+ ions. (These data were taken with our older, Q=12000 detector. Its half-width, $\gamma_0/(4\pi)$, was about 6 Hz.) In these scans, we swept the trap voltage and observed peaks when the ions' resonant frequency matched the sum of the drive frequencies. We changed one of the drive frequencies ν_d between the scans, but kept the drive amplitudes constant in the three scans.

As predicted, the resonance became narrower when the ions were excited further and further from the detector. The ions' peak amplitude should grow like γ^{-1} (Equation V.C.14), while the resonant detector becomes less sensitive to the ions' current like $\gamma^{1/2}$. Thus the amplitude of the detected current *grows* like $\gamma^{-1/2}$. Since both the widths and the peaks in Figure V.D.2 behaved in the proper way, we were confident that the widths we measured were due to the ions "natural" damping, and not, for example, to instabilities in the trapping potential or in the internal dynamics of the ion cloud. Therefore we felt some confidence in using the width of swept resonances (in general) to count larger clouds in the trap.

Anharmonicity

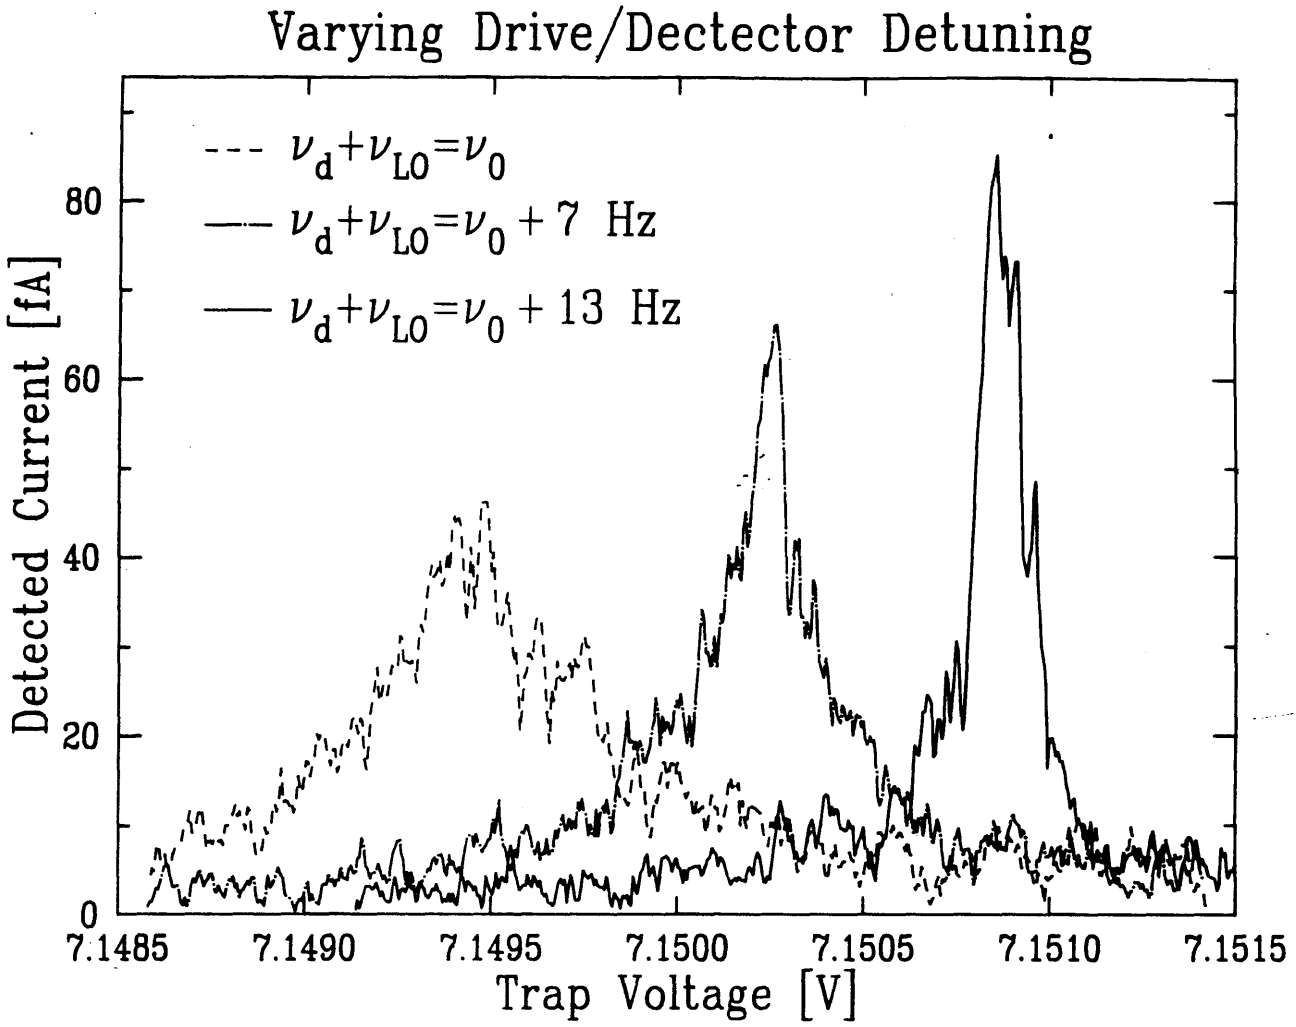

Figure V.D.2. The ions' damping (and thus their peak, resonant response) depends strongly on the detuning between the ions and the detector. The ions were moved approximately one half-width of the detector between sweeps.

In addition to ion counting, we have used the two-drive CW scheme to measure, and compensate for, trap anharmonicity. Because the two-drive scheme lets us focus on the ions' response, we can observe the anharmonic effects described in Chapter II. Although we will discuss anharmonicity and the guard rings in more detail in Chapter VII, we will illustrate here the theory of Chapter II with some examples from our experiment.

Figure V.D.3 shows a CW swept resonance of a single ion. Because the trap was not perfectly harmonic, we were able to observe hysteresis, as described in Section II.B. (See, for example, Figure II.B.2b) In Figure V.D.3, we drove the ion to about 20% of the trap. However, the anharmonicity caused an 0.6 Hz shift in the ion's resonant frequency at this amplitude (c.f Equation II.B.18). Since the natural width of this resonance is $\approx$ 0.05 Hz, the anharmonic shift is many times the natural width, and, as discussed in Equation II.B.21, we observed different behavior in the two different sweep directions. We will discuss these single ion sweeps further in Chapter VII.

By changing the potential on the guard ring we could control the amount of anharmonicity in the trap. Figure V.D.4 shows how the line-shape of a larger cloud (about 30 ions, for variety) changed when we adjusted the guard ring potential, about 50 mV across the whole figure. The detected signal in Figure V.D.4 is about the same as the previous figure because we drove the ions about 30 times less hard. Since this drive did not exceed the critical drive required to induce hysteresis, the anharmonicity

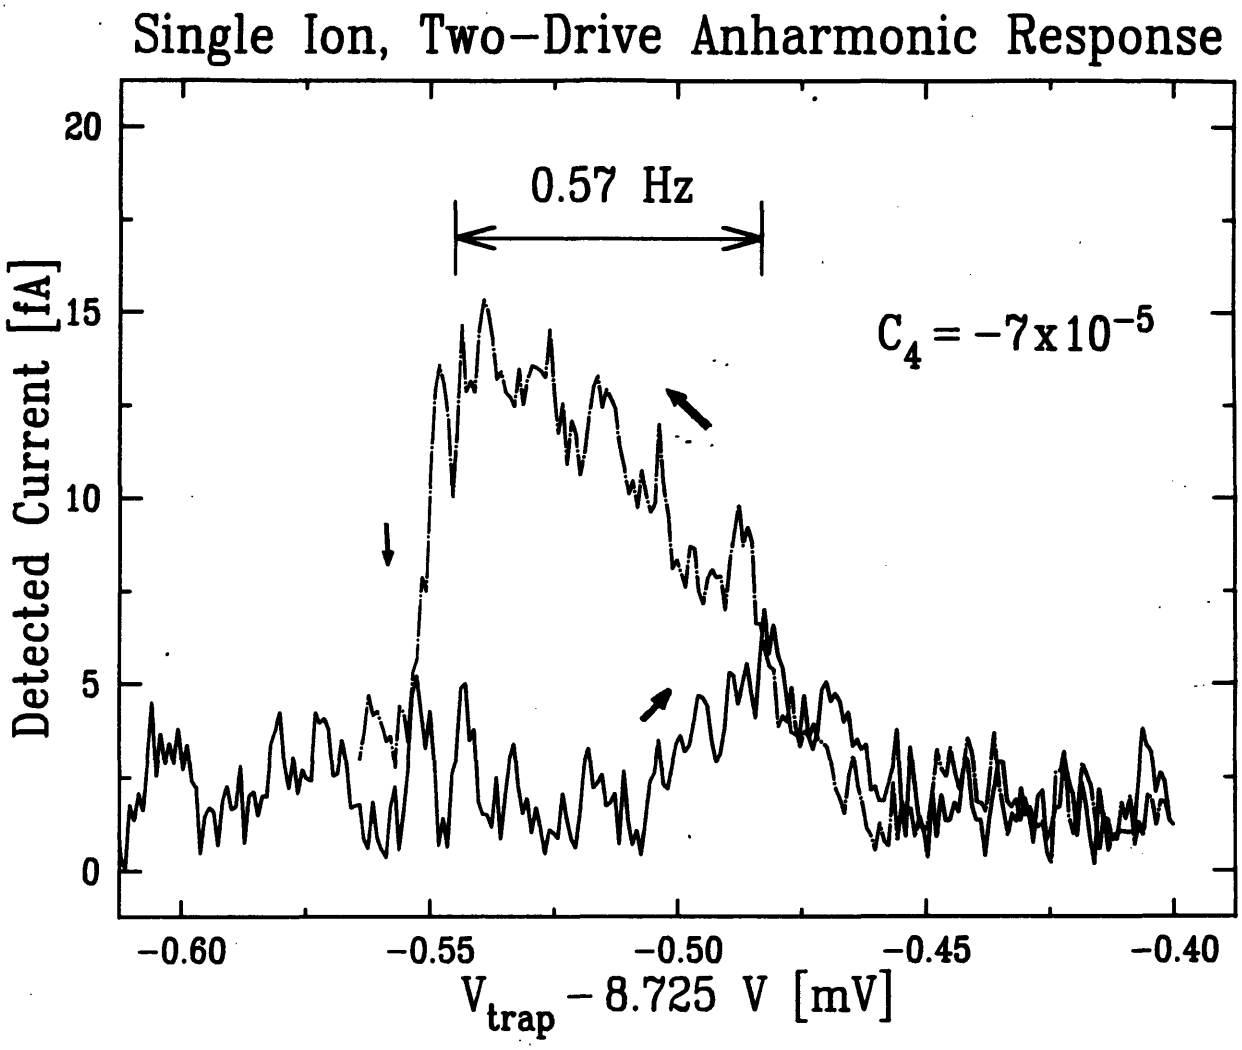

Figure V.D.3. Single ion anharmonicity. One N_2^+ ion, driven 20% of the trap size, demonstrates hysteresis.

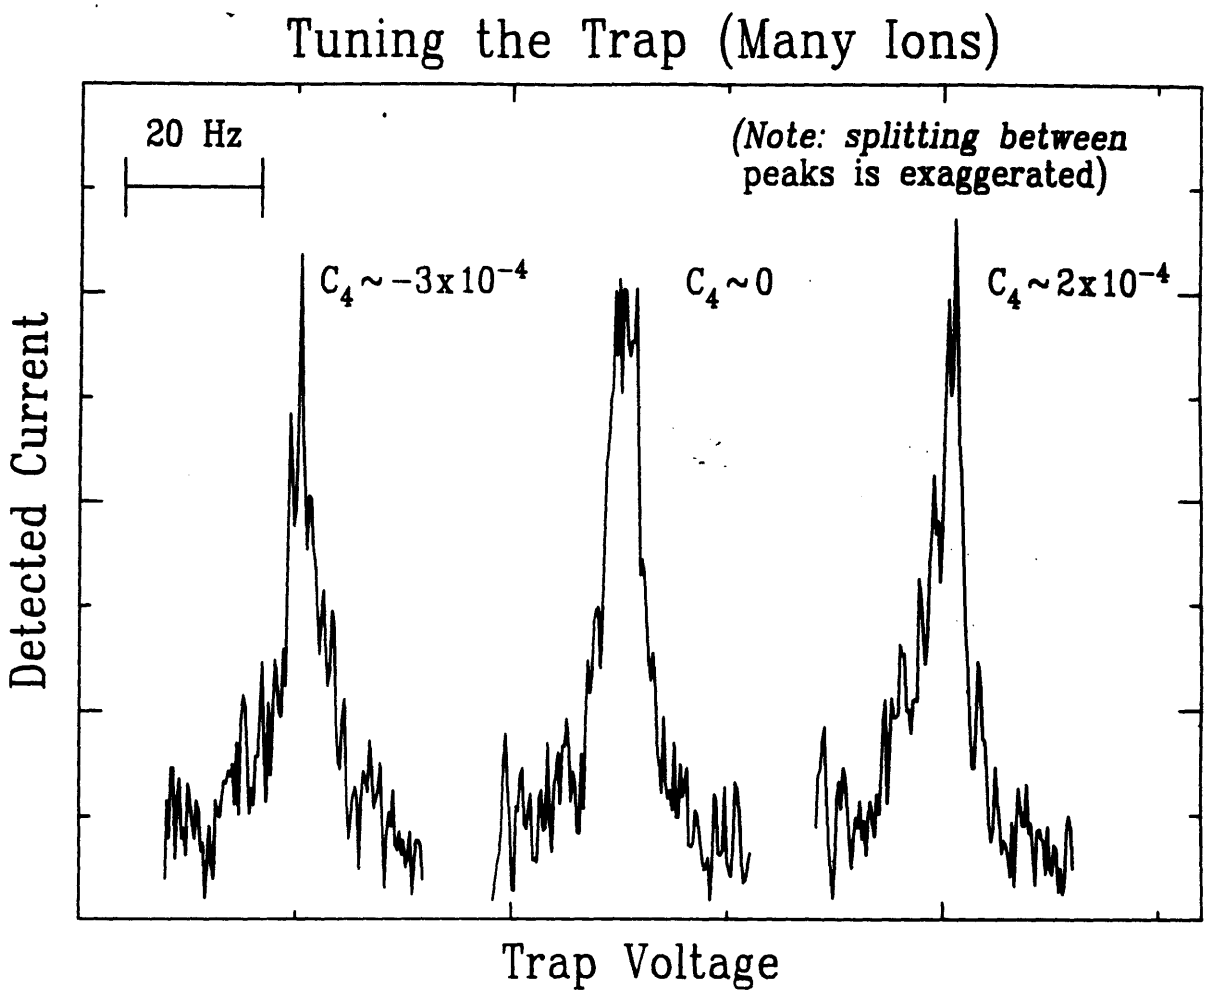

Figure V.D.4. Tuning the guard rings by symmetrizing the line-shape for clouds of ions. The sign of the asymmetry is different on the left and right line-shapes.

changed the symmetry of the peaks without causing sweep-direction dependent results.

Notice, that the many-ion cloud has more symmetric line-shapes than the single ion, Figure V.D.3, even though, all three cases, the anharmonicity is larger than the single ion case. This effect is an illustration of the idea expressed in Section II.B: for equal amplitude *output* signals, the effects of anharmonicity scale like $\gamma_z^{-3/2}$, and thus are much worse for smaller clouds.

In the presence of drift and jitter, the two-drive scheme becomes even more useful. For example, the frequency jitter of small clouds of trapped ions can be comparable to their width. In this case, the dip discussed in Section V.A cannot be observed. However, when we use the two-drive scheme to excite the ions far from their resonance, the jitter becomes inconsequential: we can adjust the detuning to be much larger than the jitter, and then increase the drives to recover the lost signal.

Finally, we have used the two-drive method to lock the ions to an external frequency reference. The feedback system was very simple. We took the orthogonal part the resonant signal (Figure V.C.3), amplified it, and added it to the D.C. trap bias. (Since both the trap bias offset and the lock-in signal were produced by the computer, this addition was very easy to implement.)

At this stage in the experiment, locking remains a curiosity, letting us tell, for example, when bad ions are shifting the axial resonant frequency (Section II.C.2), or enabling us to look for jitter in the axial frequency (we didn't have the signal-to-noise

to see any). However, as we shall discuss in Chapter VII, in order to make very high precision cyclotron measurements, we probably will need to keep the axial frequency locked. We have not yet locked a single ion—because of signal-to-noise constraints—and thus expect that we will refine our locking techniques when we attempt 10^{-10} cyclotron measurements.

V.E Two-Drive Pulsing

We have pulsed the ions using the two frequency drives. For practical reasons, we usually pulse ω_{LO} and leave on ω_d constantly. (When pulsed this way, the coil remains unexcited by the wings of the pulse since $\omega_{LO} \ll \omega_z$.) In this section, we will use the Green function, $\overline{G_z}$, to find the amplitude of the ions' response.

We have, at the end of a pulse of duration T, amplitude V_{LO} centered at frequency ω_{LO}, in the presence of a CW drive of amplitude V_d at frequency ω_d, a response:

$$z(t>T) \sim \int_0^T dt' J_0(\beta) J_1(\beta) \sin(\omega_z' t - (\omega_z' - \omega_{LO})t') V_d \cos\omega_d t' \qquad \text{(V.E.1)}$$

(I have ignored damping.) Throwing out the term oscillating at the high frequency $\omega_z' + \omega_d - \omega_{LO}$ yields:

$$z(t>T) \sim \frac{1}{2} V_d J_0(\beta) J_1(\beta) \int_0^T dt' \cos(\omega_z' t - (\omega_z' - (\omega_d + \omega_{LO}))t') \qquad \text{(V.E.2)}$$

When the pulse is shorter than the detuning, $(\omega_z - (\omega_{LO} + \omega_d))T \ll 1$, we get:

$$z(t>T) \sim \frac{1}{2} T J_0(\beta) J_1(\beta) V_d \qquad \text{(V.E.3)}$$

Once again, the two-drive scheme gives identical results as the one drive implementation except that V_d is replaced by $J_0(\beta) J_1(\beta) V_d$. Notice that this result is true even though V_d is not the drive that is pulsed! The two drives again help simplify the interpretation of the data: in these pulse experiments, the coil is not directly excited and the ions decay as if decoupled from the detector.

In Figure V.E.1 we show two-drive, pulsed results similar to the one-drive pulse experiment described in Figure V.B.2. Each peak on Figure V.E.1 represents the average of 40 transients from the 30 N_2^+ ions stored in the trap. The trap potential was changed to shift the ions about 2 Hz between each set of transients. The ions were excited to about 3% the trap size by a 10 mS pulse at $\omega_{LO} = (2\pi)1KHz$ in all the peaks. Unlike the single-drive case, the excitation amplitude does not depend on the ion/detector detuning.

However, the discrete fourier transform and the detector tuning do change the relative amplitude of the resulting peak in frequency space. That is, we again must use Equation (V.B.1) to account for the peak heights. Combining this relationship with the decreased gain off the resonance of the detector (as discussed in Section V.B, the gain falls off like $\gamma_z^{1/2}$), we see that the resulting FFT peaks should depend on the ions' damping:

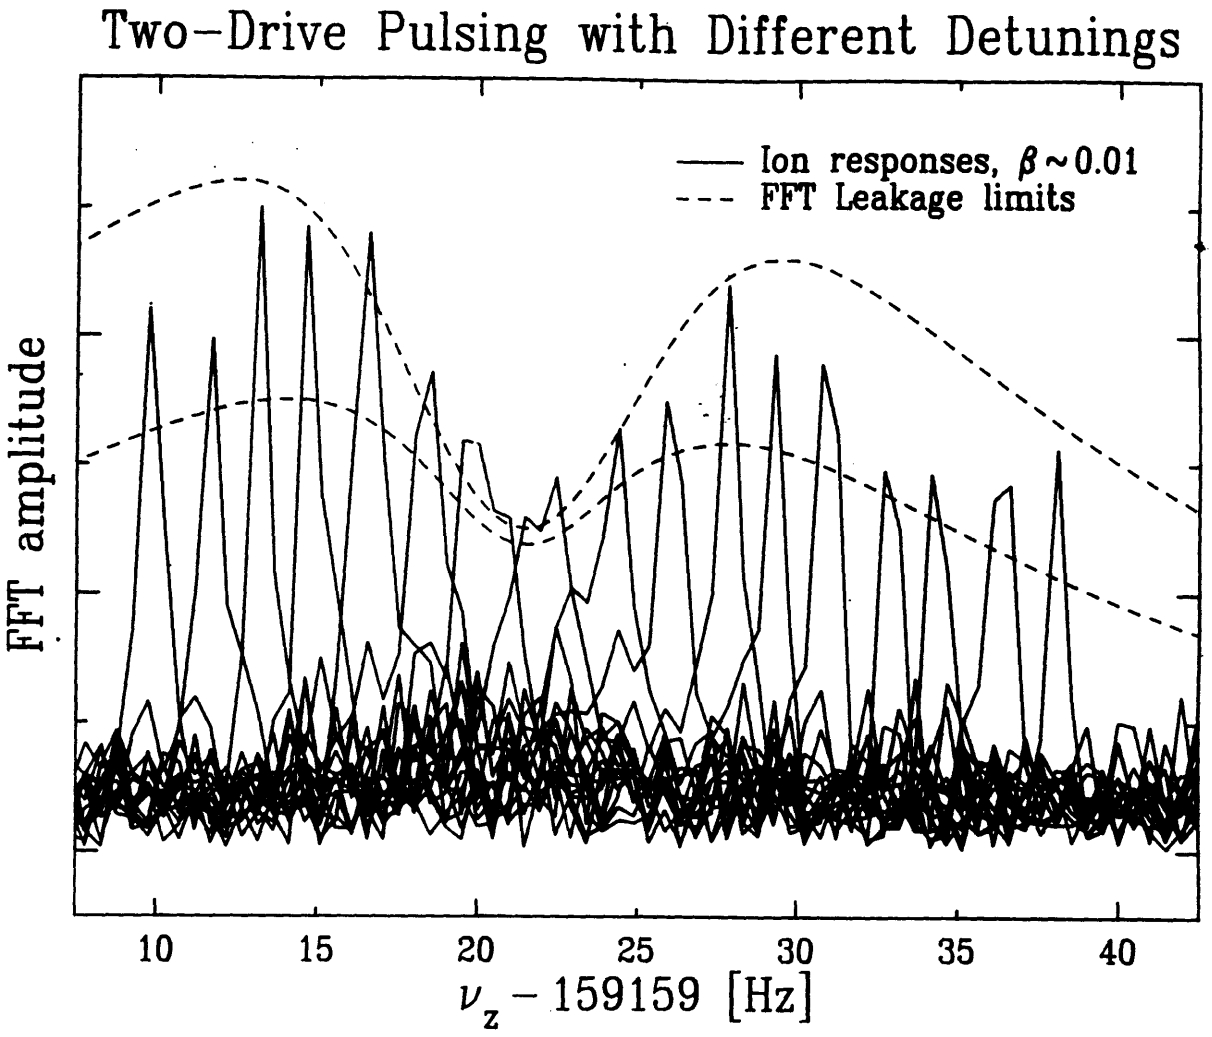

Figure V.E.1. Pulsing with two drives. Raw FFT data for lightly-driven ions, moved by 2 Hz intervals between scans. The minimum in the middle—at the location of the detector—is due to the ions' increased damping on resonance.

$$peak \sim \frac{f_{samp}}{\gamma_z^{1/2}} \left(1 - e^{-\frac{N\gamma_z}{2f_{samp}}}\right) \qquad \text{(V.E.4)}$$

Thus near the coil, we expect the peak dependence to be dominated by the $\gamma^{-1/2}$ term. Therefore when the ions are right on the detector (γ at its maximum value), the peak height should be at a local minimum. For larger detunings, $\gamma > 2f_{samp}/N$, though, the exponential kills off this dependence, and the peaks will depend on $\gamma^{1/2}$. Therefore the peak heights eventually decrease like the detuning. Thus we expect a maximum peak amplitude to occur somewhere near, but still somewhat detuned from, the detector.

We can find the damping for which Equation V.E.4 predicts a maximum peak. The resulting transcendental equation has the simple numerical solution:

$$\gamma_z^{peak} = \frac{2.51}{T} \qquad \text{(V.E.5)}$$

where T is the total length of the transient, N/f_{samp}.

The dashed lines on Figure V.E.1 show the expected limits between which the peaks should lie. (They actually represent a one parameter fit— the total amplitude was allowed to vary.) We can only establish limits because of scalloping losses (see III.B.1.13) in the FFT. That is, because the frequency of the ion signals was not known—and the signal-to-noise is insufficient to determine it to sufficient accuracy—the amplitude of the peak can be as much as 30% reduced because of leakage. (As indicated on the figure, scalloping is much less important for the broader peaks near the detector: damping already causes the peak to be spread over many bins of the

FFT.)

The asymmetry in the data (peaks are larger near $\nu_z = 0$) was a natural result caused by the duration of the pulses. The pulsed excitation is only flat (in frequency space) when the excitation-ion detuning, Δ, is much smaller than the effective (angular) frequency width of the excitation, $1/\tau$, where τ is the duration of the pulse. In general, the frequency dependence of the excitation goes like $\frac{\sin\frac{\Delta\tau}{2}}{\frac{\Delta\tau}{2}}$, (a "sinc" function). This sinc dependence has been included in the dashed curves on Figure V.E.1, and is apparent in the data.

Despite all these advantages over one-drive pulses, the two-drive scheme has a fundamental limitation. In our experiment, we cannot get enough drive voltage on the trap to excite the ions any more than was shown in Figure V.E.1. The amplitude of V_d (the high-frequency, CW drive) is limited by the dynamic range of the SQUID. Too much drive causes the SQUID to lose lock (see Section IV.A) and thus makes detection impossible. The amplitude for the pulsed (ring) drive is limited by the bessel function structure. Thus turning up this drive does increase the total excitation of the ions, but only spreads the excitation to additional sidebands.

The single-drive pulse does not suffer from this particular problem because of the detector boost. For the present detector, this boost gives about a factor of 40 increase in the effect pulse drive. Without this boost,then, the two-drive pulsed

scheme, for all its apparent advantages, will most likely remain a curiosity.

CHAPTER VI

COOLING AND DETECTING THE RADIAL MODES

We have discussed, in great detail, now, the axial motion and its detection, leaving the radial modes aside. In the this chapter, we shall consider these radial modes and describe the methods we can use to control and detect the magnetron and cyclotron motions. While other groups have split the trap ring to detect directly the cyclotron motion (the "quad-ring" method of Van Dyck, et. al., [VMF85]), we use indirect methods to detect and cool the radial modes. The main advantage of our methods is that they leaves the cyclotron motion undamped (except when we choose to damp it), and therefore the resonance lines can be made arbitrarily narrow. In addition, because we detect only the axial motion, (whose frequency can be set by changing the trapping potential) we can measure different species without changing our detector or the magnetic field.

However, to understand these methods, we must introduce a bit more formalism. In this section, we will discuss the use of an oscillating, inhomogeneous electric field to cool the magnetron and cyclotron motions. In Section VI.B, we will extend these concepts to include cyclotron resonances, detected via the axial motions. This second, more rigorous section will also contain more correct expressions for the magnetron and cyclotron cooling rates.

We want to "cool" the radial modes in order to control the resonance lines (both axial and cyclotron) for precision measurements. As discussed in Section II.C, a

whole host of perturbations scale as the square of the mode amplitudes. A lack of control over initial amplitudes therefore can translate directly into uncertainty in our final measurements. For example, though we try to create ions at the center of the trap ($\rho = 0$), some slight misalignment of the field emitter (our e^- source) is inevitable. However, as we shall see, it should be possible to *electronically* cool the magnetron radius to smaller than 10 μ, far better than any possible alignment. In addition, the ability to decrease the magnetron radius provides a cure for radial diffusion caused by impurity collisions (Section II.C.1), thus making trapping times practically indefinite.

The cyclotron motion, as we shall see, cannot be cooled nearly as effectively, though heavier mass ions ($m > 15$ amu or so) can be cooled somewhat. Even in those cases, the final cyclotron temperature still will be far greater than the background 4 K. This fact is important to the ultimate precision of the experiment: since our measurement schemes involve coupling to the detector (and its thermal noise), we actually may heat the cyclotron motion (in the thermodynamic sense) while measuring it. On the other hand, collisions between strongly-driven ions and background gas accidentally can impart a large velocity to the cyclotron motion. In such cases, then, cyclotron cooling would become essential.

VI.A Cooling the Magnetron Mode

In our trap, we use an oscillating, quadrupole electric field to couple together the axial and radial modes. We have broken the azimuthal symmetry of the trap by

splitting the guard rings (the small electrodes between the ring and endcaps whose other purpose is to compensate for the anharmonicity of the trap—see Section II.B.3). By driving the two halves of the split electrode withe the appropriate phase, we can set up an inhomogeneous field in the center of the trap. The potential varies both along the z axis and the axis of the split, which we call "x" for convenience.

To see how this field can couple the axial and radial modes, imagine an ion moving in a magnetron orbit in the z=0 plane. (Ignore any cyclotron motion.) Because the potential goes like zx, the axial force depends linearly on x and the radial force depends on z. (See Figure VI.A.1) Assume that the coupling drive, at frequency ω_d, is in phase with the ion at point A on Figure VI.A.1. The ion thus receives an upward kick at $t = 0$. We want the coupling drive to kick the ion again, in phase with the first kick, when the ion comes around its magnetron orbit again at a time $t = 2\pi/\omega_m$ later. That is, at this later time we want an integer multiple of 2π radians to have passed in both the coupling drive and the axial motion. Therefore, the difference between those frequencies must satisfy $(\omega_d - \omega_z)t = 2\pi n$ where n is any integer. Using t above, we see that:

$$\omega_d = \omega_z + n\,\omega_m \qquad \text{(VI.A.1)}$$

The sign of n has a profound impact on the results of the coupling drive. As we shall see, for $n = +1$ the coupling drive tends to equilibrate the magnetron and axial motions. For $n = -1$, on the other hand, the coupling tends to increase both the magnetron and axial motions exponentially, eventually driving the ion out of the trap.

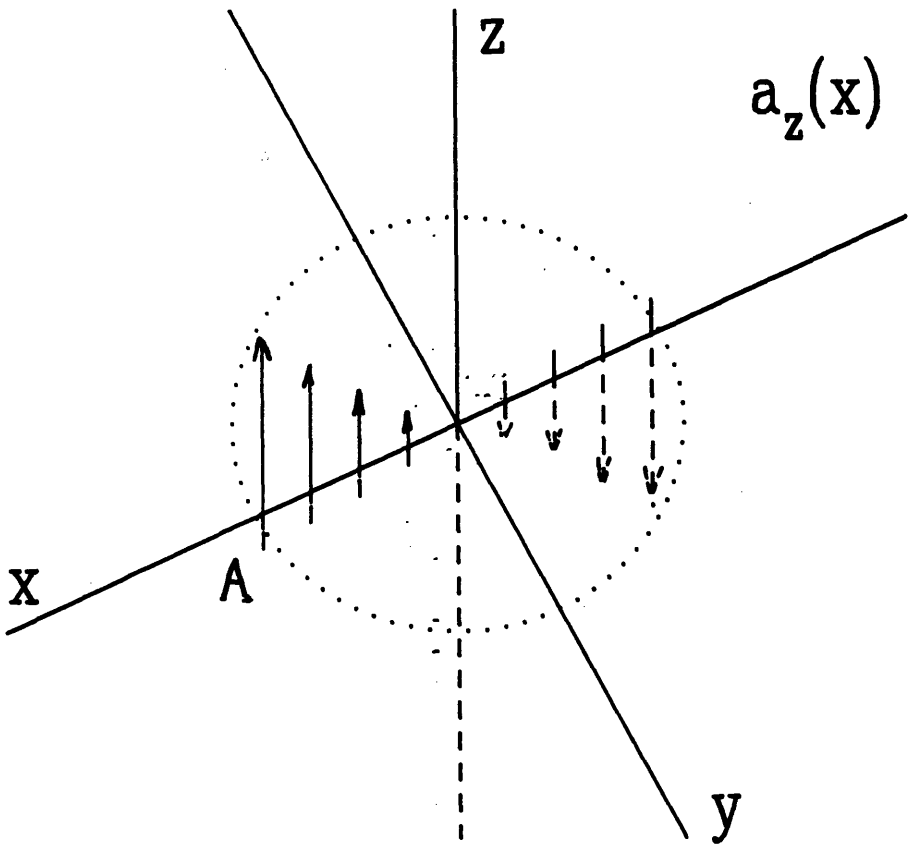

Figure VI.A.1. An inhomogeneous electric field causes a force in the z-direction which depends on the x-position.

In order to understand the difference between heating and cooling, we now return to the ion's equations of motion. When driven by a voltage, $V_d \cos\omega_d t$, the split guard ring creates a potential, near the center of the trap:

$$\Phi = \xi V_d \frac{xz}{d^2} \cos\omega_d t \tag{VI.A.2}$$

where ξ is a geometric shielding factor. The acceleration in the x direction, a_x, is thus:

$$a_x = -\frac{e\xi V_d}{md^2} z \cos\omega_d t \tag{VI.A.3}$$

$$= -az\,\omega_z^2 \cos\omega_d t$$

where:

$$a \equiv \xi \frac{V_d}{V_{trap}} \tag{VI.A.4}$$

where a is a dimensionless constant which much be measured empirically. Inserting this acceleration into the radial equation of motion, (II.A.2.6), we find:

$$\ddot{\vec{\rho}} - \omega_c\, \hat{z} \times \dot{\vec{\rho}} - \frac{1}{2}\omega_z^2 \vec{\rho} = \begin{bmatrix} -az\,\omega_z^2 \cos\omega_d t \\ 0 \end{bmatrix} \tag{VI.A.5}$$

Changing to the $\vec{V}^{(\pm)}$ variables, and explicitly writing out the vector components, we can express (VI.A.5) as:

$$\begin{bmatrix} \dot{V}_x^{\pm} \\ \dot{V}_y^{\pm} \end{bmatrix} - \omega_{\mp} \begin{bmatrix} -V_y^{\pm} \\ V_x^{\pm} \end{bmatrix} = \begin{bmatrix} -az\,\omega_z^2 \cos\omega_d t \\ 0 \end{bmatrix} \tag{VI.A.6}$$

Focusing our attention on the magnetron motion (the lower signs in VI.A.6), we find:

$$\ddot{V}_y^{(-)} - \omega_-^2 V_y^{(-)} = -az\,\omega_-\omega_z^2\cos\omega_d t \qquad \text{(VI.A.7)}$$

The acceleration in the z direction, from (VI.A.2), is:

$$a_z = -ax\,\omega_z^2\cos\omega_d t \qquad \text{(VI.A.8)}$$

We can find x in terms of the $\vec{V}^{(\pm)}$ using Equation (II.A.2.10). Taking the difference between $\vec{V}^{(+)}$ and $\vec{V}^{(-)}$, we find:

$$x = \frac{V_y^{(+)} - V_y^{(-)}}{\omega_+ - \omega_-} \qquad \text{(VI.A.9)}$$

and thus we can close the system of coupled equations between $V_y^{(-)}$ and z with:

$$\ddot{z} + \gamma_z\,\dot{z} + \omega_z^2 z = \frac{a\,\omega_z^2}{\omega_+ - \omega_-} V_y^{(-)}\cos\omega_d t \qquad \text{(VI.A.10)}$$

We have ignored the cyclotron part of x ($V_y^{(+)}$) because it is oscillating at an unimportant frequency. In addition, we have ignored the trap capacitance and assumed that the only effect of the detector is to give the ions some small damping, γ_z. (We will relax these last conditions in the next section.)

We can use these equations of motion to demonstrate that coupling at $\omega_d \approx \omega_z + \omega_-$ will cool the magnetron and $\omega_d \approx \omega_z - \omega_-$ will heat it up. We can also estimate the rates for these processes, deferring until the next section a more precise calculation. Consider the cooling or heating due a coupling drive exactly resonant between the two modes, $\omega_d = \omega_z \pm \omega_-$. Let us begin with some magnetron motion, $V_y^{(-)} = A\,e^{i\omega_- t}$. From (VI.A.7), we see that the component of $V_y^{(-)}\cos\omega_d t$ oscillating near ω_z, $\frac{1}{2}V_y^{(-)}e^{\mp i\omega_z t}$, will drive the axial mode strongly, resulting in motion

$z = \pm\frac{i}{2}\frac{a\omega_z}{\gamma_z(\omega_+ - \omega_-)}Ae^{\mp i\omega_z t}$. Putting this motion back into the radial equation, (VI.A.7) and again ignoring the non-resonant term yields, for the added radial force:

$$a_x = \mp\frac{ia^2\omega_-\omega_z^3}{4\gamma_z(\omega_+ - \omega_-)}A\,e^{i\omega t} \qquad \text{(VI.A.11)}$$

$$= \mp\frac{a^2\omega_z^3}{4\gamma_z(\omega_+ - \omega_-)}\dot{V}_y^{(-)}$$

Thus, depending on the sign of this term, the coupling either provides stable damping ($a_x \sim -\dot{V}_y^{(-)}$) or unstable heating ($a_x \sim +\dot{V}_y^{(-)}$). With the coupling drive at $\omega_d = \omega_z \pm \omega_-$ we thus have an effective "damping" constant:

$$\gamma_{mag}^{\pm} = \pm\frac{a^2\omega_z^3}{4\gamma_z(\omega_+ - \omega_-)} \qquad \text{(VI.A.12)}$$

Therefore, the upper sideband truly damps the radial motion, whereas the lower sideband leads to exponential growth the both the radial and axial motions. As we shall see in Section VI.B, this expression is optimistic: it is only true for sufficiently weak coupling drives, such that $\gamma_{mag}^{+} \ll \frac{1}{4}\gamma_z$. For stronger drives, the cooling saturates at $\gamma_z/2$, though the heating continues unbounded.

The discussion above applies equally to cyclotron-axial coupling at $\omega_d = \omega_+ \pm \omega_z$. In that case, however, we see from (VI.A.9) that the axial acceleration has the opposite sign, and therefore the role of the sidebands is reversed: $\omega_+ - \omega_z$ cools while $\omega_+ + \omega_z$ heats.

Our language has been a little sloppy here. When we write "heating" and "cooling," we do not mean them in the strict thermodynamic sense. Rather, by heating we only mean an exponential growth of the orbit; by cooling, exponential damping. In fact, for the magnetron motion, we do not even mean heating or cooling in the energy sense at all! Because of the sign of the radial electrostatic potential, we "cool" the magnetron motion (that is, decrease its orbit) by *increasing* its energy. To keep down confusion, then, we choose colloquial meanings; cooling always means reducing the orbit sizes, heating always means increasing them.

Figure VI.A.2 shows the results of a typical magnetron cooling cycle. Right after loading about 20 N_2^+ ions (and attempting to rid the trap of other ions), we set one of our HP3325 oscillators to sweep its frequency continuously, back and forth between 162341 and 162357 Hz, at about 120 seconds per sweep. (This was ±8 Hz around the expected magnetron cooling frequency, $\nu_z + \nu_-$.) There was no axial drive. Each time the cooling drive swept through the cooling frequency, we observed axial action being transferred from the magnetron motion by the cooling drive. We recorded the axial noise with a $\tau = 0.25$ s filter, centered at the axial frequency, 159603 Hz. (We used the lock-in program described in Section IV.E.3.)

In the lower graph, we plotted the oscillator frequency as a function of time and marked with an "x" the approximate location of the axial noise peaks. They all fall at about 162351(1) Hz, indicating a magnetron frequency of 2748 Hz. We typically saw no additional cooling after about 500 seconds, indicating that the cooling had reached

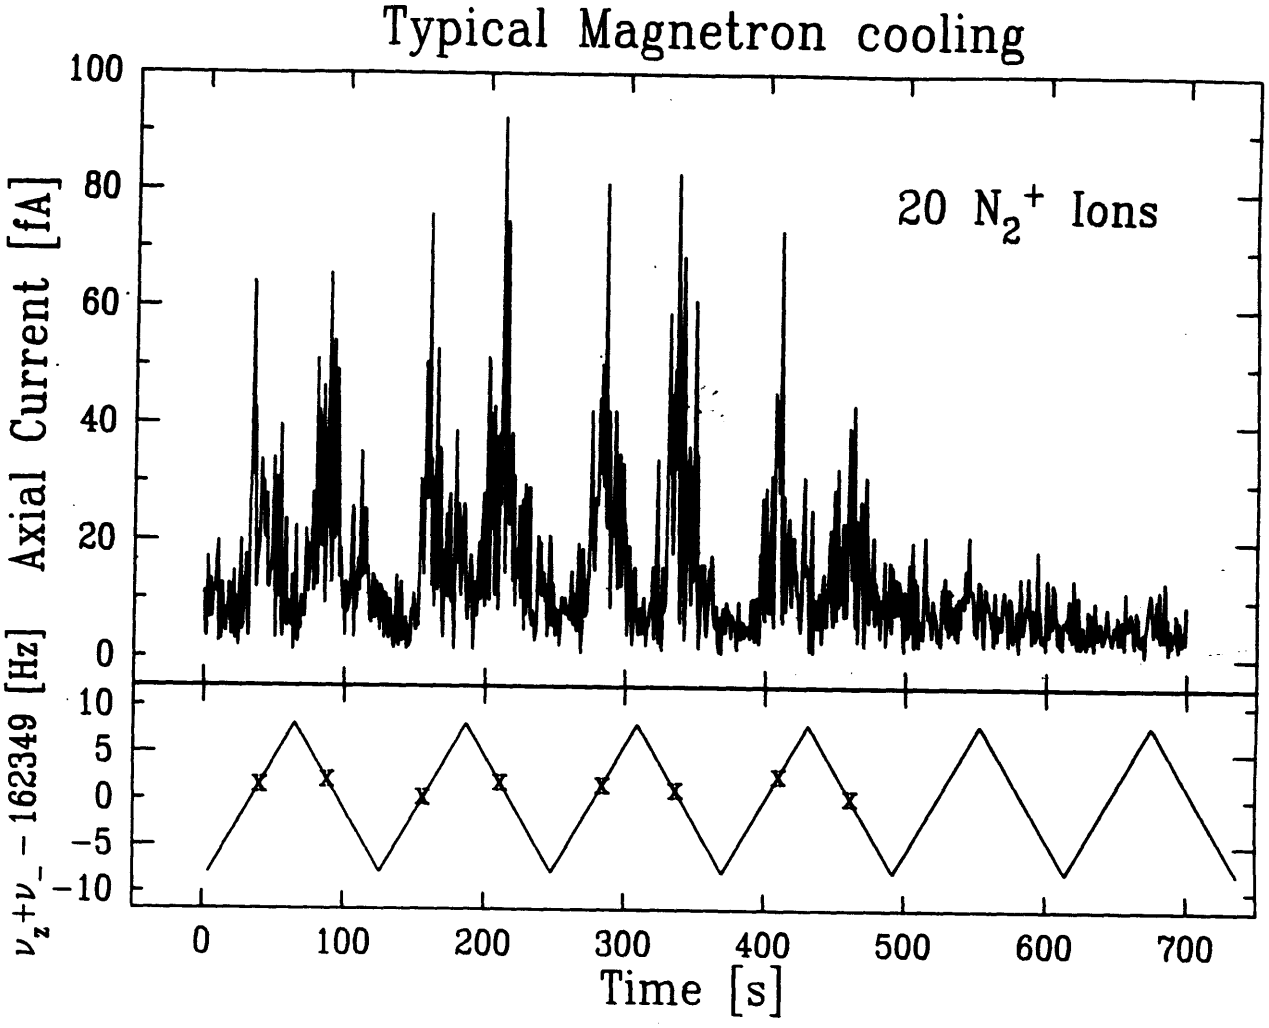

Figure VI.A.2. Magnetron cooling. The upper graph shows the axial noise (as a function of time) when the magnetron/axial coupling frequency was swept across the cooling resonance. The lower curve shows the coupling frequency as a function of time.

some limit.

The peak induced axial signals were quite large. If they came from a coherent axial motion, its equivalent amplitude would be about 10% of the trap size. Using the equal-action argument to be presented in the next sub-section, this axial amplitude could come from coupling to a magnetron radius of about 0.15 mm, consistent with a reasonable misalignment between the field-emitter and the electrical center of the trap.

Cooling Limit

This cooling proceeds until it reaches a limit, set by thermodynamics. The axial detector provides a heat bath which bring the axial motion into thermal equilibrium at some temperature, T_z, the effective temperature of the amplifier. Though the cooling limit can be given by quantum mechanical arguments [BRG86], we derive the limit here using purely classical ideas. The core of the argument is that a quadrupole RF coupling, linear in both coordinates, conserves the sum of the action of the two modes.

For a harmonic oscillator of mass m, displacement y, using the usual conserved quantities, we find that the action (classically, $\int \vec{p} \cdot d\vec{q}$) is proportional to the energy divided by the frequency, ω:

$$S \sim \frac{E}{\omega} = \frac{1}{2}\frac{m}{\omega}\dot{y}^2 = \frac{1}{2} m \omega y^2 \qquad \text{(VI.A.13)}$$

Using the equations of motion, (VI.A.7) and (VI.A.10), we can write down the time-

derivative of the total action of the coupled axial and magnetron modes:

$$\dot{S} \sim \left\{-\dot{V}_y^{(-)} z\, \omega_z + V_y^{(-)} \dot{z}\, \omega_-\right\} \cos\omega_d t \qquad \text{(VI.A.14)}$$

where, to get the magnetron action, we have used $x \sim V_y^{(-)}/\omega_+$. For the coupling frequencies $\omega_d = \omega_z \pm \omega_-$, and the short-term mode behaviors, $Vy \sim e^{+i\omega_- t}$ and $z \sim e^{\pm i\omega_z t}$, we find:

$$\dot{S} \sim \left\{-i\,\omega_-\omega_z \pm i\,\omega_z\,\omega_-\right\} V_y^{(-)} z \qquad \text{(VI.A.15)}$$

and thus for magnetron cooling, (the upper sign), $\dot{S} = 0$. Therefore, since the total action is conserved, we can decrease the magnetron motion only by adding axial action, and vice versa. Since the axial mode has some finite temperature, it thus must have some average, minimum action. Therefore cooling cannot continue after the action in the two modes has been equilibrated to this thermal value. Thus we find a thermal limit:

$$\rho_m^2 = 2\frac{\omega_-}{\omega_z} z_{rms}^2 \qquad \text{(VI.A.16)}$$

For N_2^+ ions, using a detector at 4 K, this thermal limit becomes $z_{rms} = 3.5\times10^{-3}$ cm and $\rho_m = 6.5\times10^{-4}$ cm. Another way of expressing this limit is as an effective temperature for the magnetron mode [BRG86]:

$$T_m = -\frac{\omega_-}{\omega_z} T_z \qquad \text{(VI.A.17)}$$

where the negative sign expresses the fact that the magnetron motion is inherently

unstable. Since $\omega_- \gg \omega_z$, we therefore see that the magnetron motion will be cooled far below the axial temperature.

This limit also applies to the cyclotron cooling (on the proper sideband, $\omega_d = \omega_+ - \omega_z$). However, since the cyclotron frequency is higher than the axial frequency, the cyclotron motion will equilibrate at a temperature much greater than the detector temperature:

$$T_c = \frac{\omega_+}{\omega_z} T_z \qquad \text{(VI.A.18)}$$

Thus, cyclotron "cooling" can actually produce greater random fluctuation in the cyclotron motion if, for example that motion had been cooled near T_z by collisions with cold background gas. In terms of the the thermal axial motion, the cyclotron cooling limit can be expressed as an effective cyclotron radius, ρ_c:

$$\rho_c^2 = \frac{\omega_z}{\omega_+} z_{rms}^2 \qquad \text{(VI.A.19)}$$

Since $2\omega_-\omega_+ = \omega_z^2$, we see that the magnetron and cyclotron cooling limits produce identical orbit sizes. Thus, for N_2^+ at 4K, the cyclotron limit also yields $\rho_c = 6.5\times10^{-4}$ cm.

VI.B Detecting the Cyclotron Mode

We have developed two techniques to measure the cyclotron frequency of our trapped ions. In the first, we measure axial frequency shifts induced by a very strong cyclotron cooling drive. In the second, we drive the cyclotron mode directly and measure the axial motion coherently induced by the cooling drive. The first method, very simple in practice, is rather complicated to describe theoretically. The second method, rather more difficult to implement, is a simpler to explain. In this section, I will describe the framework we use to explain both resonant techniques.

This framework is an extension of the operator method used in [BRG86], extended to included the effects of our resonant detector. Within this framework, we will predict both the shift in the axial resonance due to the coupling drive ("cyclotron resonance by avoided crossing") and axial motion induced because of a joint cyclotron and cooling drive ("cyclotron induced axial resonance"). As a side benefit, we will get the true, "saturable" magnetron (and cyclotron) cooling rate, alluded to in the previous section.

Let us write down the equations for the axial motion, $z(t)$, the radial motion, $V_y^{\pm}(t)$ (cyclotron or magnetron), and the voltage induced across the detector, $\mathbf{V}_1(t)$. We need to include this voltage for the same reason that we did in Section V.A: the dynamics of the detector and the ions are intricately related. For generality, let us also include, in addition to the coupling drive, $a\cos\omega_d t$, an axial drive $\mathbf{V}_{in}(t)$, and a radial

drive $V_x(t)$. The axial acceleration due to the first drive as we saw in (V.A.1.1) is $a_z = \frac{eB_1}{2mz_0} V_{in}(t)$. Similarly, the radial acceleration due to the second drive is $a_x = \frac{e\chi}{2mz_0} V_x(t)$, where χ is another geometric constant, here $\ll 1$. (In our experiment, we also use the split guard rings to create this radial drive. The relative phase of the two guard rings determines whether the drive is a coupling drive or a radial drive. This implementation will be discussed further in Chapter V.)

We can extend the harmonic oscillator equations (VI.A.7) and (VI.A.10) to include these axial and radial forces. The detector voltage (recall Figure V.A.2.2) also behaves like a harmonic oscillator, driven both by $V_{in}(t)$ and by the current induced by the ions' motion, (II.A.3.4):

$$\ddot{V}_1 + \frac{1}{R(C+C_t)}\dot{V}_1 + \frac{1}{L(C+C_t)}V_1 = \frac{C_t}{(C+C_t)}\ddot{V}_{in} + \frac{eB_1}{2z_0}\frac{1}{C+C_t}\ddot{z} \qquad \text{(VI.B.1)}$$

where L, R, and C are given in Figure V.A.2.2, and C_t is C_{trap} on that figure.

Using these equations we can write, in matrix form, the coupled dynamics for the ion and the detector:

$$\begin{bmatrix} \frac{d^2}{dt^2}+\omega_z^2 & 1 & \pm\frac{a\omega_z^2}{\omega_+-\omega_-}\cos\omega_d t \\ -\frac{C+C_t}{C_t}\gamma_0\gamma_z & \frac{C+C_t}{C_t}\left(\frac{d^2}{dt^2}+\gamma_0\frac{d}{dt}+\omega_0^2\right) & 0 \\ a\omega_\pm\omega_z^2\cos\omega_d t & 0 & \frac{d^2}{dt^2}+\omega_\pm^2 \end{bmatrix} \begin{bmatrix} z \\ \frac{eB_1}{2mz_0}V_1 \\ V_y^\pm \end{bmatrix} = \begin{bmatrix} a_z \\ a_z \\ a_x \end{bmatrix} \qquad \text{(VI.B.2)}$$

I have removed the damping from Equation VI.A.10 here because, in fact, that

damping comes from the detector, and thus is a result from, not an input to (VI.B.2). The frequencies and dampings, γ_0, γ_z, ω_0, and ω_z are given in (V.A.2.7) and (V.A.2.8). The matrix on the left side of (VI.B.2) is an operator, and thus the time, t (as in, $\cos\omega_d t$) is also an operator. We seek to invert this operator matrix; that is, we want a set of operators that, given the drives $a_x(t)$ and $a_z(t)$, yield the motions:

$$\begin{bmatrix} z(t) \\ \dfrac{eB_1}{2mz_0}V_1(t) \\ V_y^{\pm}(t) \end{bmatrix} = \begin{bmatrix} G_{zz} & G_{zx} \\ G_{1z} & G_{1x} \\ G_{xz} & G_{xx} \end{bmatrix} \begin{bmatrix} a_z \\ a_x \end{bmatrix} \qquad \text{(VI.B.3)}$$

This operator formalism, used in Brown and Gabrielse, is very compact. For example, the $z(t)$ term with the operator notation and time dependences spelled out says:

$$z(t) = \int_{-\infty}^{+\infty} dt' G_{zz}(t,t')a_z(t') + \int_{-\infty}^{+\infty} dt' G_{zx}(t,t')a_x(t') \qquad \text{(VI.B.4)}$$

The explicit time-dependence of the matrix (VI.B.2) destroys the time-shift invariance of the system; that is, $G_{zz}(t,t') \neq G_{zz}(t-t')$. However, as we shall see below, by neglecting some sideband structure, we can in some cases restore the shift-invariance and thus consider the fourier transform of all the variables, essentially returning to the admittance formalism of Section V.A, but now incorporating the coupled degrees of freedom.

In analogy to quantum mechanics, G_{zz} (for example) is an operator (analogous to the forward propagator [GOT79]) and $G_{zz}(t,t')$ is a matrix element of that operator, expressed $(t\,|\,G_{zz}\,|\,t')$. The constant operator yields a delta function,

$(t \mid \mathbf{1} \mid t') = \delta(t - t')$, and we have a completeness condition, $\mathbf{1} = \int dt \mid t)(t \mid$. To illustrate the use of these operators, let us make contact with our "usual" notion of Green functions. The Green function, $G(t-t')$, for the differential equation $\left[\frac{d^2}{dt^2} + \omega_z^2\right] z = a_z(t)$ must satisfy [JAK83]:

$$\left[\frac{d^2}{dt^2} + \omega_z^2\right] G(t,t') = \delta(t - t') \tag{VI.B.5}$$

We consider $G(t,t')$ and $\delta(t - t')$ as the matrix elements of the operators G and $\mathbf{1}$, respectively. Let us introduce another function, $G^{-1}(t,t')$ that is given by:

$$G^{-1}(t,t') = \left[\frac{d^2}{dt^2} + \omega_z^2\right] \delta(t-t') \tag{VI.B.6}$$

We can re-write (VI.B.5) by inserting a complete set of states:

$$\left[\frac{d^2}{dt^2} + \omega_z^2\right] \int dt_A \; \delta(t - t_A) G(t_A,t') = \delta(t - t') \tag{VI.B.7}$$

Expressing this relation as a matrix element, we have:

$$\int dt_A \; (t \mid G^{-1} \mid t_A)(t_A \mid G \mid t') = (t \mid \mathbf{1} \mid t') \tag{VI.B.8}$$

and thus, integrating over the states $\mid t_A)$:

$$G^{-1}G = \mathbf{1} \tag{VI.B.9}$$

I have belaboured this explanation to show that the actual manipulation of these operators becomes straight forward algebra when we remember what the notation means.

Now, defining:

$$G_z^{-1} \equiv \frac{d^2}{dt^2} + \omega_z^2 \qquad \text{(VI.B.10a)}$$

$$G_1^{-1} \equiv \frac{C + C_t}{C_t} \iint \left[\frac{d^2}{dt^2} + \gamma_0 \frac{d}{dt} + \omega_0^2 \right] \qquad \text{(VI.B.10b)}$$

$$G\pm_{-1} \equiv \frac{d^2}{dt^2} + \omega_\pm^2 \qquad \text{(VI.B.10c)}$$

we can re-write (VI.B.2) and (VI.B.3) as:

$$\begin{bmatrix} G_z^{-1} & 1 & \pm\frac{a\,\omega_z^2}{\omega_+ - \omega_-}\cos\omega_d t \\ -\frac{C + C_t}{C_t}\gamma_0\gamma_z & G_1^{-1} & 0 \\ a\,\omega_\pm\omega_z^2\cos\omega_d t & 0 & G_\pm^{-1} \end{bmatrix} \begin{bmatrix} G_{zz} & G_{zx} \\ G_{1z} & G_{1x} \\ G_{xz} & G_{xx} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad \text{(VI.B.11)}$$

We thus have six equations for the six, unknown operators in the middle matrix. Different operators will be important for different applications. For the avoided crossing measurement, we are interested in G_{1z}: this operator produces the detector voltage from an axial drive in the presence of the cooling drive. On the other hand, for the induced resonance, G_{zx} produces the axial ion response from the cyclotron drive. Yet another operator, G_{xx} will contain the information that determines the true magnetron damping.

Obviously, I won't give all the steps in solving these equations, just the highlights. To solve, we first write the mixing operators (G_{zx}, G_{1z}, G_{1x}, G_{xz}) in terms of the "diagonal" operators, G_{zz} and G_{xx}:

$$G_{xz} = -a\,\omega_{\pm}\omega_z^2 G_{\pm}\cos\omega_d t\, G_{zz} \qquad \text{(VI.B.12a)}$$

$$G_{1z} = G_1\left[1+\frac{C+C_t}{C_t}\gamma_0\gamma_z\, G_{zz}\right] \qquad \text{(VI.B.12b)}$$

$$G_{zx} = \pm\frac{a\,\omega_z^2}{\omega_+-\omega_-}\left[G_z^{-1}+\frac{C+C_t}{C_t}\gamma_0\gamma_z\, G_1\right]^{-1}\cos\omega_d t\, G_{xx} \qquad \text{(VI.B.12c)}$$

$$G_{1x} = \mp\frac{C+C_t}{C_t}\frac{a\,\omega_z^2\gamma_0\gamma_z}{\omega_+-\omega_-}G_1\left[G_z^{-1}+\frac{C+C_t}{C_t}\gamma_0\gamma_z\, G_1\right]^{-1}\cos\omega_d t\, G_{xx} \qquad \text{(VI.B.12d)}$$

Using the equation, we can solve for G_{zz}:

$$G_{zz} = \left\{G_z^{-1}+\frac{C+C_t}{C_t}\gamma_0\gamma_z G_1 \mp \frac{a^2\omega_z^4\omega_{\pm}}{\omega_+-\omega_-}\cos\omega_d t\, G_{\pm}\cos\omega_d t\right\}^{-1}\left[1-G_1\right] \qquad \text{(VI.B.13)}$$

and for G_{xx};

$$G_{xx} = \left\{G_{\pm}^{-1} \mp \frac{a^2\omega_{\pm}\omega_z^4}{\omega_+-\omega_-}\cos\omega_d t\left[G_z^{-1}+\frac{C+C_t}{C_t}\gamma_0\gamma_z G_1\right]^{-1}\cos\omega_d t\right\}^{-1} \qquad \text{(VI.B.14)}$$

To use these expressions, we must evaluate terms like $\cos\omega_d t\, G\cos\omega_d t$. Following Brown and Gabrielse, we can insert several complete sets of states and write:

$$(t\,|\cos\omega_d t\, G\cos\omega_d t\,|\,t') = \int dt_A\, dt_B\,(t\,|\cos\omega_d t\,|\,t_A)(t_A\,|\,G\,|\,t_B)(tB\,|\cos\omega_d t\,|\,t') \qquad \text{(VI.B.15)}$$

$$= G(t,t')\left\{\frac{1}{2}\cos\omega_d(t-t') + \frac{1}{2}\cos\omega_d(t+t')\right\}$$

The $\cos\omega_d(t+t')$ term causes sidebands and, unlike similar terms in Section V.C, these sidebands are always off-resonance. We therefore disregard them. Thus, when $G(t,t')$ is shift-invariant, we have returned the shift-invariance to the problem. In that case, we can simplify everything by using fourier transforms. Using tildes to mark

fourier transforms, and using the sign convention:

$$\tilde{G}(\omega) = \int dt\, G(t) e^{-i\omega t} \qquad \text{(VI.B.16)}$$

(the *opposite* sign convention of BRG86), we find the fourier transform of the right-hand side of (VI.B.15) = $\frac{1}{4}\tilde{G}(\omega+\omega_d) + \frac{1}{4}\tilde{G}(\omega-\omega_d)$.

We now have sufficient tools to evaluate $\tilde{G}_{zz}$ in the presence of cyclotron cooling and heating. To use (VI.B.13), we must evaluate $\tilde{G}_+(\omega\pm\omega_d)$ when $\omega \approx \omega_z$ and $\omega_d \approx \omega_+ \pm \omega_z$. (The top sign is heating, the bottom sign is cooling.) Now $\tilde{G}_+(\omega)$ is only large near $\pm\omega_+$, and thus only one of two sidebands will contribute. Defining the coupling detuning, Δ_d, by $\omega_d = \omega_+ \pm \omega_z + \Delta_d$, we find, for the $\cos\omega_d t\, G_+ \cos\omega_d t$ term in G_{zz}:

$$\frac{1}{4}\left[\tilde{G}_+(\omega+\omega_d) + \tilde{G}_+(\omega-\omega_d)\right] \approx \pm\frac{1}{8\omega_+(\omega-\omega_z \mp\Delta_d)} \qquad \text{(VI.B.17)}$$

Therefore, using:

$$1 - \tilde{G}_1(\omega) \approx \frac{C}{C+C_t}\frac{\omega^2 - i\gamma_0\omega - \omega_{LC}^2}{\omega^2 - i\gamma_0\omega - \omega_0^2} \qquad \text{(VI.B.18)}$$

(where, as before, $\omega_{LC} = (LC)^{-1/2}$) we can write, in the presence of a cyclotron coupling drive:

$$\tilde{G}_{zz}^{(+)}(\omega) = -\frac{C}{C+C_t}\frac{\omega^2 - i\gamma_0\omega - \omega_{LC}^2}{(\omega^2-\omega_z^2)(\omega^2 - i\gamma_0\omega - \omega_0^2) - \gamma_0\gamma_z\omega^2 \pm \dfrac{a^2\omega_z^4}{8(\omega_+-\omega_-)}\dfrac{\omega^2 - i\gamma_0\omega - \omega_0^2}{\omega-\omega_z+\Delta_d}}$$

When $a = 0$ (coupling drive off), we recover the admittance equation (V.A.2.23). For $a \neq 0$, we can simplify (VI.B.19) using the notation of Section V.A:

$$\tilde{G}_{zz}^{(+)}(\omega) = \frac{-\Delta_{LC}}{\Delta_z(\Delta_0 - i\gamma_0/2) - \frac{1}{4}\gamma_0\gamma_z \pm \frac{a^2\omega_z^3}{16(\omega_+ - \omega_-)}\frac{\Delta_0 - i\gamma_0/2}{\Delta_z + \Delta_d}} \qquad \text{(VI.B.20)}$$

We can use this expression together with (VI.B.12b) to find the detector admittance, $\tilde{G}_{1z}^{+}(\omega)$:

$$\tilde{G}_{1z}(\omega) \approx \frac{C}{C_t}\frac{\omega_0}{2}\frac{\Delta_z \pm \frac{a^2\omega_z^3}{16(\omega_+ - \omega_-)}\frac{1}{\Delta_z + \Delta_d}}{\Delta_z(\Delta_0 - i\gamma_0/2) - \gamma_0\gamma_z/4 \pm \frac{a^2\omega_z^3}{16(\omega_+ - \omega_-)}\frac{\Delta_0 - i\gamma_0/2}{\Delta_z + \Delta_d}} \qquad \text{(VI.B.21)}$$

Compare this equation to its zero-coupling analogue, (V.A.2.6). In that earlier equation (and here when $a = 0$) the ions short out the coil and produce a single dip in the frequency response at $\Delta_z = 0$. However, when $a \neq 0$, we see from (VI.B.21) that we can get *two* dips because the numerator can have two zeros. The location of these dips are the solutions to the quadratic equation:

$$\Delta_z^2 + \Delta_z\Delta_d \pm \frac{a^2\omega_z^3}{16(\omega_+ - \omega_-)} = 0 \qquad \text{(VI.B.22)}$$

And thus:

$$\Delta_z = \frac{-\Delta_d \pm \sqrt{\Delta_d^2 \mp \frac{a^2\omega_z^3}{4(\omega_+ - \omega_-)}}}{2} \qquad \text{(VI.B.23)}$$

My notation here is a bit awkward. The first $\pm$ refers to the two roots of the quadratic equations, but the second $\mp$ inside the radical refers to the heating/cooling sidebands in ω_d. For the heating sideband (the minus sign in the radical), we see that a strong heating drive does not lead to a zero because Δ_z is complex. For a cyclotron cooling

drive, on the other hand, both roots are real, and we therefore expect two dips in the frequency response. The separation of these dips depends characteristically on Δ_d.

When the cooling drive is exactly resonant between the two states ($\Delta_d = 0$), we see that the two dips are separated by a frequency Δ_{avoid}:

$$\Delta_{avoid} \equiv \sqrt{\frac{a^2\omega_z^3}{4(\omega_+ - \omega_-)}} \qquad \text{(VI.B.24)}$$

which is directly proportional to the strength of the coupling, a. The form of the axial frequency shift is identical to the standard avoided level-crossing formula [CDL77] for the energy shift of a two-level system coupling by an oscillating interaction. Exploiting this analogy, we call these shifts of the axial frequency a "classical avoided crossing." When the frequency of the perturbation equals the energy splitting between the two levels, the states repel each other, avoiding by an amount proportional to the strength of the interaction.

In Figure VI.B.1, we show a plot of the axial dips versus the detuning of the cooling drive. As discussed in Section V.A, when we excite the ions by pulsing the axial drive (V_{in}), the admittance (in this case, $\tilde{G}_{1z}(\omega)$) gives the frequency response of the detector. We can therefore measure both the cyclotron frequency and the strength of the cooling drives by exciting the axial motion with a pulse $V_1(\omega)$ and recording the response for different cooling drive frequencies. When the cooling drive is nearly resonant, $\omega_d \approx \omega_+ - \omega_z$, we will observe the double dip in the response instead of the usual single minimum. We can then plot the frequencies of the dips from different

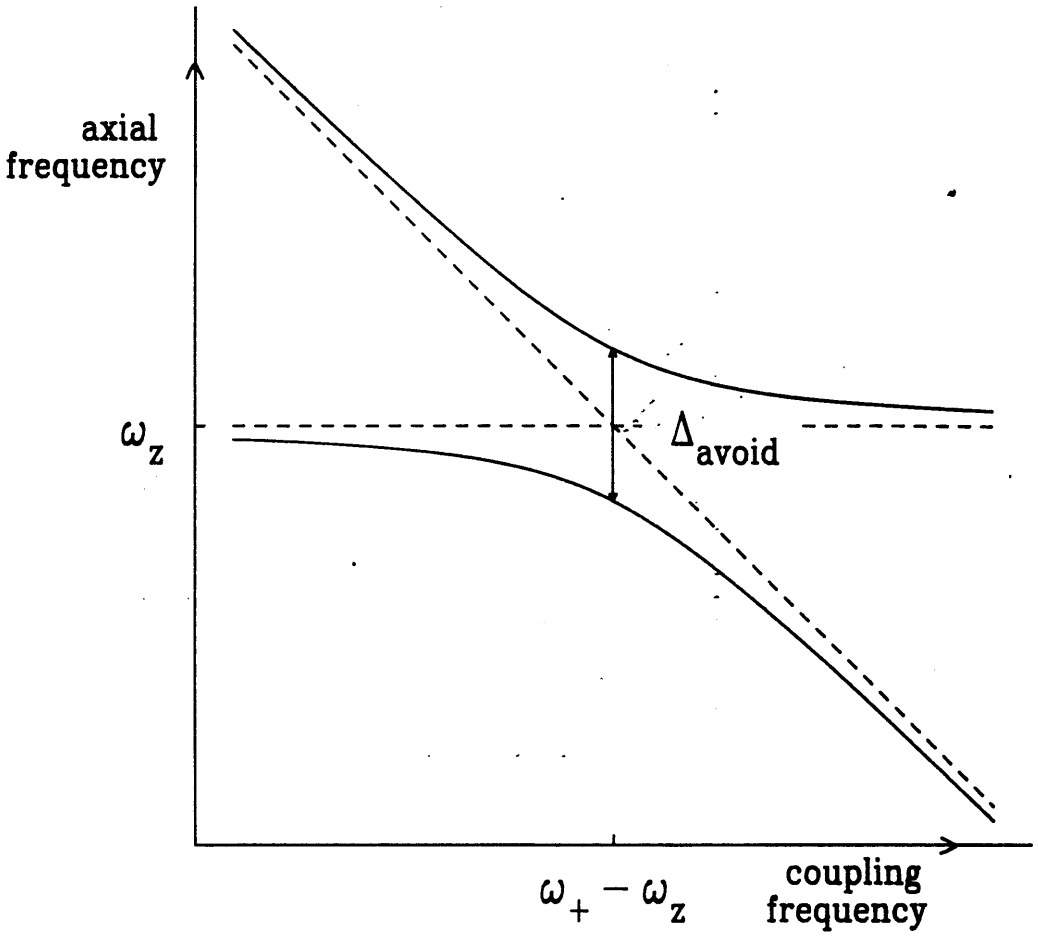

Figure VI.B.1. "Avoided crossing" in the axial resonance frequency as a function of the axial/cyclotron coupling drive frequency. Right on the coupling resonance, the two axial resonances are separated by Δ_{avoid} which is directly proportional to the strength of the coupling drive.

cooling drive frequencies against coupling detunings, Δ_d, and fit the avoided-crossing line-shape, Equation (VI.B.23) through the dips.

In order to understand the true axial resonances, we can re-write the denominator of (VI.B.20):

$$\left[\Delta_z \pm \frac{a^2\omega_z^3}{16(\omega_+-\omega_-)}\frac{1}{\Delta_z+\Delta_d}\right]\left[\Delta_0 - i\gamma_0/2\right] - \frac{1}{4}\gamma_0\gamma_z \qquad \text{(VI.B.25)}$$

We see that the resonant structure of the axial motion is identical to that for the uncoupled case, (V.A.2.6), except that the coupling drive causes Δ_z to be replaced by $\Delta_z \pm \frac{a^2\omega_z^3}{16(\omega_+-\omega_-)}\frac{1}{\Delta_z+\Delta_d}$. Thus we see that an effect of the cooling drive is to shift the axial resonance ω_z. Since the axial damping is frequency dependent, the real part of the detector's impedance at the shifted ω_z will determine the axial damping, (V.A.2.12).

We have performed cyclotron measurements using this method. One avoided crossing is presented in Figure VI.B.2. The data on this graph are the measured axial frequency of the dips against the cyclotron coupling drive frequency. These data come from eleven, one-drive excitations of 15 N_2^+ ions to 5% of the trap size. In four of these excitations, both dips were clearly visible in the FFT spectra. We then fit the frequency of these dips to the avoided-crossing line shape, Equation VI.B.23, letting the cyclotron frequency and the strength of the coupling, a, vary.

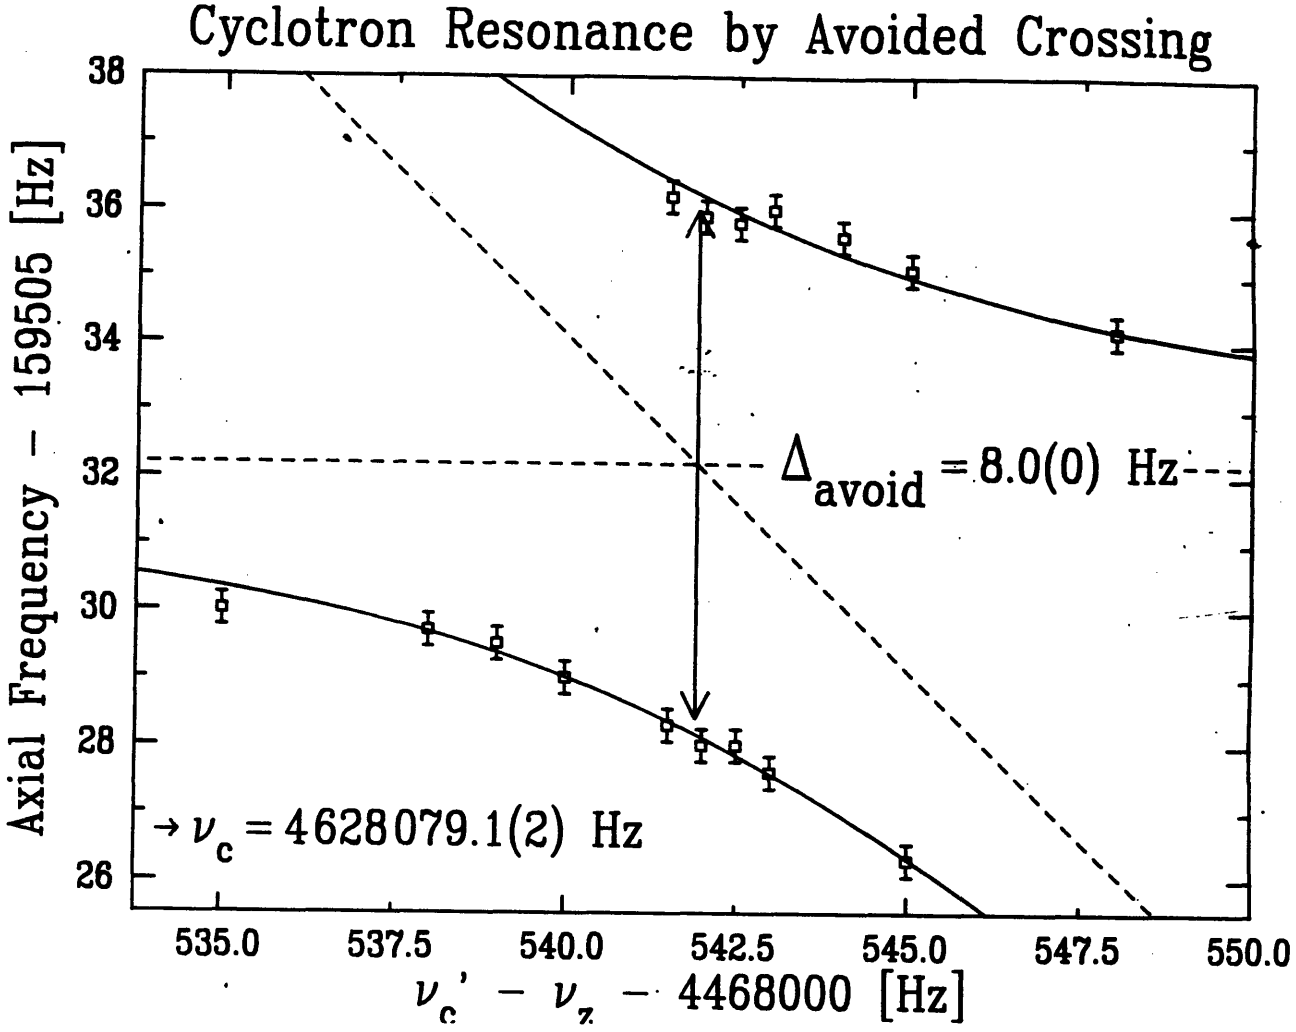

Figure VI.B.2. An actual cyclotron measurement by avoided crossing. The N_2^+ ions were pulsed with a one-drive scheme and the frequencies of the resulting "dips" were recorded as a function of the coupling frequency.

From this fit, we determined the cyclotron cooling frequency, $\nu_c' - \nu_z = 4468541.9(2)$ Hz. Adding on the (independently) measured axial frequency, $\nu_z = 159537.2(1)$ Hz, we obtain the cyclotron frequency in the trap, 4628079.1(2). Using the quadratic sum rule, (II.C.7.4), we get a final, free-space cyclotron value, $\nu_c = 4630828.8(2)$ Hz. However, before claiming these as 50 ppb measurements of the true cyclotron frequency, we should note that the errors above are the statistical and do not include systematic effects.

In fact, the possible systematic effects point to some of the immediate drawbacks of using avoided crossings for precision measurements. Since these measurements were made with the single-drive pulse scheme, we know from Section V.A that the resulting ion amplitude will depend on the detuning of the ion from the coil. However, the cyclotron coupling changes this detuning, and thus changes the amplitudes of the axial responses. These different amplitudes will cause different frequency shifts due to the various perturbations described in Section II.C. For example, if our MACOR trap does indeed create a 6 gauss$/cm^2$ magnetic bottle, then the 5% excitation used in this experiment should cause roughly 0.5 Hz shift in the cyclotron frequency! This shift would be several times our statistical error, and would be different at different coupling frequencies.

In addition, in this scheme, we must add the axial frequency to the measured coupling frequency to obtain ν_c'. Therefore, errors in the axial frequency become directly errors in the measured cyclotron frequency. Since we will *never* know the

axial frequency to ultra-high precisions of 10^{-11}, (or even to the 10^{-9} level), there is no reason to extend this technique beyond its present accuracy.

Though the avoided-crossing scheme can't provide high-precision data, it does give us a measurement of the cyclotron coupling strength, a. Since our best a priori guess for this value depends strongly on shielding factors which are very difficult to estimate, this coupling strength calibration was particularly useful.

True Cyclotron/Magnetron Damping Constant

We can determine the effective damping of the radial modes by evaluating $\tilde{G}_{xx}(\omega)$. As we might imagine from Equation VI.B.14, $\tilde{G}_{xx}(\omega)$ has a very similar structure to $\tilde{G}_{zz}(\omega)$. We must evaluate an expression like (VI.B.15), but this time for an operator G_2:

$$G_2 = \left[G_z^{-1} + \frac{C + C_t}{C_t} \gamma_0 Gz\, G_1 \right]^{-1} \qquad \text{(VI.B.26)}$$

This Green operator is beginning to look familiar (c.f. V.A.2.23). It is time-invariant, and its fourier transform is:

$$\tilde{G}_2(\omega) = -\frac{\omega^2 - i\gamma_0\omega - \omega_0^2}{(\omega^2 - \omega_z^2)(\omega^2 - i\gamma_0\omega - \omega_0^2) - \omega^2\gamma_0\gamma_z} \qquad \text{(VI.B.27)}$$

To use the prescription (VI.B.15), we need to evaluate $G_2(\omega \pm \omega_d)$ for $\omega \approx \omega_+$, again with $\omega_d = \omega_+ \pm \omega_z + \Delta_d$:

(VI.B.28)

$$\frac{1}{4}\left[G_2(\omega+\omega_d)+G_2(\omega-\omega_d)\right]=\pm\frac{1}{8\omega_z}\frac{\omega-\omega_+-\Delta_d\mp(\omega_z-\omega_0\mp i\gamma_0/2)}{(\omega-\omega_+-\Delta_d)(\omega-\omega_+-\Delta_d\mp(\omega_z-\omega_0\mp i\gamma_0/2)-\gamma_0\gamma_z/4}$$

Plugging back into (VI.B.14) we get an expression for $\tilde{G}_{xx}(\omega)$:

(VI.B.29)

$$\tilde{G}_{xx}(\omega)=\frac{1}{2\omega_+}\frac{-1}{\omega-\omega_+\pm\frac{a^2\omega z^3}{16(\omega_+-\omega_-)}\left\{\frac{\omega-\omega_+-(\Delta_d\pm(\omega_z-\omega_0\mp i\gamma_0/2}{(\omega-\omega_+-\Delta_d)(\omega-\omega_+-(\Delta_d\pm(\omega_z-\omega_0\mp i\gamma_0/2))}\right\}}$$

As usual, the poles of $\tilde{G}_{xx}(\omega)$ give the resonances: the real part yields the resonant frequency; the imaginary part, half the damping. In the most general case, though, we see that we must solve a cubic equation to finds these poles. However, in the case $\gamma_z \ll \gamma_0/4$ (the weakly-coupled case from Section V.A), we can use (V.A.10-12) to simplify the term in the braces:

$$\tilde{G}_{xx}(\omega)=\frac{-1}{2\omega_+}\frac{1}{\omega-\omega_+\pm\frac{a^2\omega_z^3}{16(\omega_+-\omega_-)}\frac{1}{\omega-\omega_+-\Delta_d-i\gamma_{eff}/2}} \qquad \text{(VI.B.30)}$$

where γ_{eff} is the effective damping which comes from the real part of the detector impedance at ω_z, given by Equation (V.A.2.12). We can solve the quadratic equation for the root of (VI.B.30) that, as $a \to 0$ connects to $\omega = \omega_+$:

$$\omega-\omega_+=\frac{\Delta_d+i\gamma_{eff}/2-\sqrt{(\Delta_d+i\gamma_{eff}/2)^2\mp\frac{a^2\omega_z^3}{4(\omega_+-\omega_-)}}}{2} \qquad \text{(VI.B.31)}$$

In the limit of small a, we see that $\omega_d = \omega_+ - \omega_z$ (the lower sideband) cools, and the cooling rate is proportional to the square of the coupling. However, for very strong cooling drive, $\frac{a^2\omega_z^3}{4(\omega_+-\omega_-)} \gg \gamma_{eff}^2$, we see that the radical will be positive for all Δ_d,

and thus the entire damping will be given by $i\gamma_{eff}/4$. Since the axial damping saturates at half the coil damping, we find that the cyclotron damping must saturate at one quarter the coil damping.

For magnetron coupling, $\tilde{G}_{xx}(\omega)$ is almost identical, except that the role of the sidebands is reversed, and the upper sideband cools. For small coupling drive, we recover our earlier approximation, (VI.A.12). For large coupling, the magnetron damping, too, saturates at $\gamma_{eff}/2$.

Before leaving this topic, note that there is also a small shift in the cyclotron resonance because of the coupling, given by the real part of the pole, (VI.B.31). (This is the complementary effect to the axial shift discussed above.) See further the paragraph above Equation V.A.2.12 for more discussion.

Induced Resonance

Building on our discussion of G_{xx} and G_{zz}, we can use the operator formalism to determine the axial response to an additional *cyclotron* drive, a_x. Unlike the axial response to an axial drive, the detector and ion response to a cyclotron drive are very similar. We can see from (VI.B.12b and c) that:

$$G_{1x} = \frac{C + C_t}{C_t} \gamma_0 \gamma_z G_1 G_{zx} \qquad \text{(VI.B.32)}$$

and thus the only difference between G_{1x} and G_{zx} is the resonance enhancement of

the induced current because of the high-Q detector (the G_1 term.) That is, when there is no direct axial drive, we can think of the detector simply as a current amplifier, effectively increasing the induced current by the Q of the detector, ω_0/γ_0.

We see from (VI.B.12c) that we have a term in G_{zx} that looks like:

$$G = G_A \cos\omega_d t \, G_B \tag{VI.B.33}$$

Evaluating the matrix element of such a term, $G(t,t')$ gives:

$$G(t,t') = \int \frac{d\omega}{2\pi} \tilde{G}_A(\omega) \left\{ \frac{1}{2} \tilde{G}_B(\omega-\omega_d) e^{+i\omega_d t'} + \frac{1}{2} \tilde{G}_B(\omega+\omega_d) e^{-i\omega_d t'} \right\} e^{i\omega(t-t')} \tag{VI.B.34}$$

Notice that this green function depends not only on $(t-t')$, but also has t' dependence. Therefore, we cannot explicitly fourier transform an operator like G_{zx}. However, we can determine the coherent response to a harmonic drive. For example, we can determine the action of our (hypothetical) operator $G(t,t')$ on a drive $e^{i\omega_A t'}$. Calling this response $r(t)$:

$$r(t) = \int_{-\infty}^{+\infty} dt' \, G(t,t') e^{i\omega_A t'} \tag{VI.B.35}$$

Plugging in (VI.B.34), doing the dt' integral first, then integrating $d\omega$ over the resulting delta-function we see that:

$$r(t) = \frac{1}{2} \tilde{G}_B(\omega_A) \left[\tilde{G}_A(\omega_A+\omega_d) e^{i(\omega_A+\omega_d)t} + \tilde{G}_A(\omega_A-\omega_d) e^{-i(\omega_A-\omega_d)t} \right] \tag{VI.B.36}$$

Thus operators like G produce sidebands at $\pm\omega_d$.

Combining this result with our calculation of $\tilde{G}_{xx}(\omega)$, we can compute the amplitude of these sidebands for G_{zx}, the ion's axial response to the cyclotron drive.

We have from (VI.B.12c) and (VI.B.26):

$$G_{zx} = -\frac{a\omega_z^2}{\omega_+ - \omega_-} G_2 \cos\omega_d t \, G_{xx} \qquad \text{(VI.B.37)}$$

We can use the result (VI.B.36) above, writing $\omega_A = \omega_+ + \Delta_x$, and $\omega_d = (\omega_+ + \Delta_x) - (\omega_z + \Delta_z)$. We consider simultaneously sweeping the cyclotron drive and the cooling drive, with their frequency difference held constant at $\omega_z + \Delta_z$. We then find:

$$z(t) \approx \frac{a\omega_z^2}{2(\omega_+ - \omega_-)} \tilde{G}_{xx}(\omega_+ + \Delta_x)\tilde{G}_2(\omega_z + \Delta_z) e^{i(\omega_z + \Delta_z)t} \qquad \text{(VI.B.38)}$$

and, for the radial motion,

$$V_y^{(+)}(t) \approx \tilde{G}_{xx}(\omega_+ + \Delta_x) e^{i(\omega_+ + \Delta_x)t} \qquad \text{(VI.B.39)}$$

In this method of resonance, we sweep Δ_x with Δ_z fixed. Thus, while Δ_z sets the maximum response (through $\tilde{G}_2$), the axial resonance is *observed* through the peak in $\tilde{G}_{xx}$ when Δ_x goes through zero. We can use these equations to estimate the amplitude of the axial response. When the ions are weakly coupled to the coil, we have:

$$\tilde{G}_2(\omega_z + \Delta_z) \approx \frac{-1}{2\omega_z(\Delta_z - i\gamma_{eff}/2)} \qquad \text{(VI.B.40)}$$

and thus, the full axial response is:

$$z(t) = \frac{a\omega_z}{8\omega_+^2} \frac{1}{\Delta_z - i\gamma_{eff}/2} \left\{ \frac{1}{\Delta_x - \dfrac{a^2\omega_z^3}{16(\omega_+ - \omega_-)} \dfrac{1}{\Delta_z - i\gamma_{eff}/2}} \right\} a_x \qquad \text{(VI.B.41)}$$

This expression lets us predict the width and amplitude of the cyclotron-induced axial

resonance. We can adjust the width by adjusting Δ_z: by increasing Δ_z, we decrease the imaginary part of the term in the braces in two ways. First, in the usual way, the imaginary part for large Δ_z will be given by $\frac{\gamma_{eff}}{\Delta_z^2}$. Additionally, γ_{eff} itself depends on Δ_z: detuning the ion from the coil decreases the real part of the impedance the ion sees from the detector and thus decreases γ_{eff}.

Notice that, for $\Delta_z = 0$, Equation (VI.B.41) predicts an effective width for the induced resonance:

$$\delta_{cyc}^{FWHM} \approx \frac{1}{4} \frac{a^2 \omega_z^3}{\gamma_{eff} (\omega_+ - \omega_-)} \qquad \text{(VI.B.42)}$$

This width continues to increase with increased strength of the coupling drive. Thus, like the CW one-drive resonances (recall Figure V.A.3), the width of this resonance does not saturate even though the true damping does.

We can use (VI.B.41) to compute the maximum axial response on resonance ($\Delta_x = 0$):

$$z_{peak} \approx \frac{2}{a \omega_+ \omega_z^2} \frac{\gamma_{eff}/2}{\Delta_z - i\gamma_{eff}/2} \qquad \text{(VI.B.43)}$$

At the same detuning, (VI.B.39) gives the maximum cyclotron response:

$$V_{y,peak}^{+} \approx \frac{4\gamma_{eff}}{a^2 \omega_z^3} a_x \qquad \text{(VI.B.44)}$$

Since $\rho_c \approx \frac{V_y^{(+)}}{\omega_+}$, we find that:

$$z_{peak} \approx \left\{ \frac{1}{4} \frac{a\omega_z}{\Delta_z - i\gamma_{eff}/2} \right\} \rho_c \qquad \text{(VI.B.45)}$$

Therefore the amplitude of the cooling drive changes the relative amount of cyclotron and axial motion induced on resonance. This effect can be valuable for minimizing shifts due to perturbations (because of ρ_c) without sacrificing undue signal-to-noise, which depends on z_0..

We can see the vivid differences between this induced resonance and the avoided-crossing measurements. Comparing (VI.B.24) and (VI.B.40), we see that the width of the avoided crossing and the width of the induced resonance are related:

$$\delta_{cyc}^{FWHM} = \frac{\Delta_{avoid}^2}{\gamma_{eff}} \qquad \text{(VI.B.46)}$$

Therefore, if we were to use the same cooling drive that gives a good avoided crossing, we would get an extremely broad induced cyclotron resonance. For example, if the same cooling amplitude that produces an 8 Hz avoided crossing (Figure VI.B.3) were used to drive a single N_2^+ ion's induced resonance, that resonance would be over 1 KHz wide!

Another interesting observation about the second method is the relative efficiency of the induced resonance in terms of cyclotron amplitude. If we turn down the cooling drive to the amplitude at which the true cyclotron damping just begins to saturate, we find:

$$\rho_{c,peak} = \sqrt{\frac{\omega_-}{\omega_z}}\, z_{peak} \qquad \text{(VI.B.47)}$$

That is, the relative cyclotron-to-axial response is the same as the thermodynamic cooling limit, (VI.A.19). For stronger cooling drives, there will be less radial motion for a given axial signal but a broader response. Thus we can trade the width of the resonance for smaller cyclotron orbits when we are very worried about perturbations.

It is worth reiterating here the comment from the end of Section II.C.5. In the presence of magnetic field inhomogeneities, this induced cyclotron resonance almost certainly will be hysteretic. Recall from our discussion of non-linear resonances (Section II.B) that the hysteretic regime begins when the frequency shift (due to the orbit size) equals the damping of the oscillator. Because the cyclotron damping saturates at $\gamma_{eff}/2$ (presently, 0.15 s^{-1}), small axial motion will cause shifts in the cyclotron frequency of this order. For example, using the prediction for the cyclotron shift, (II.C.5.5), and our estimated bottle, 6 gcm^{-2}, axial orbits of 0.1 mm should send us into the hysteretic regime. With our present signal-to-noise we often drive the ion ten times harder than this to detect it. Thus we understand the need to eliminate the magnetic inhomogeneities for precision cyclotron work.

Figure VI.B.3 demonstrates some of the typical difficulties encountered when we attempted these induced resonances. To obtain these data, we swept (discontinuously) two, locked HP3325 oscillators across the cyclotron resonance. One oscillator was set near the cyclotron resonance frequency, ν_c', the other near the cooling frequency,

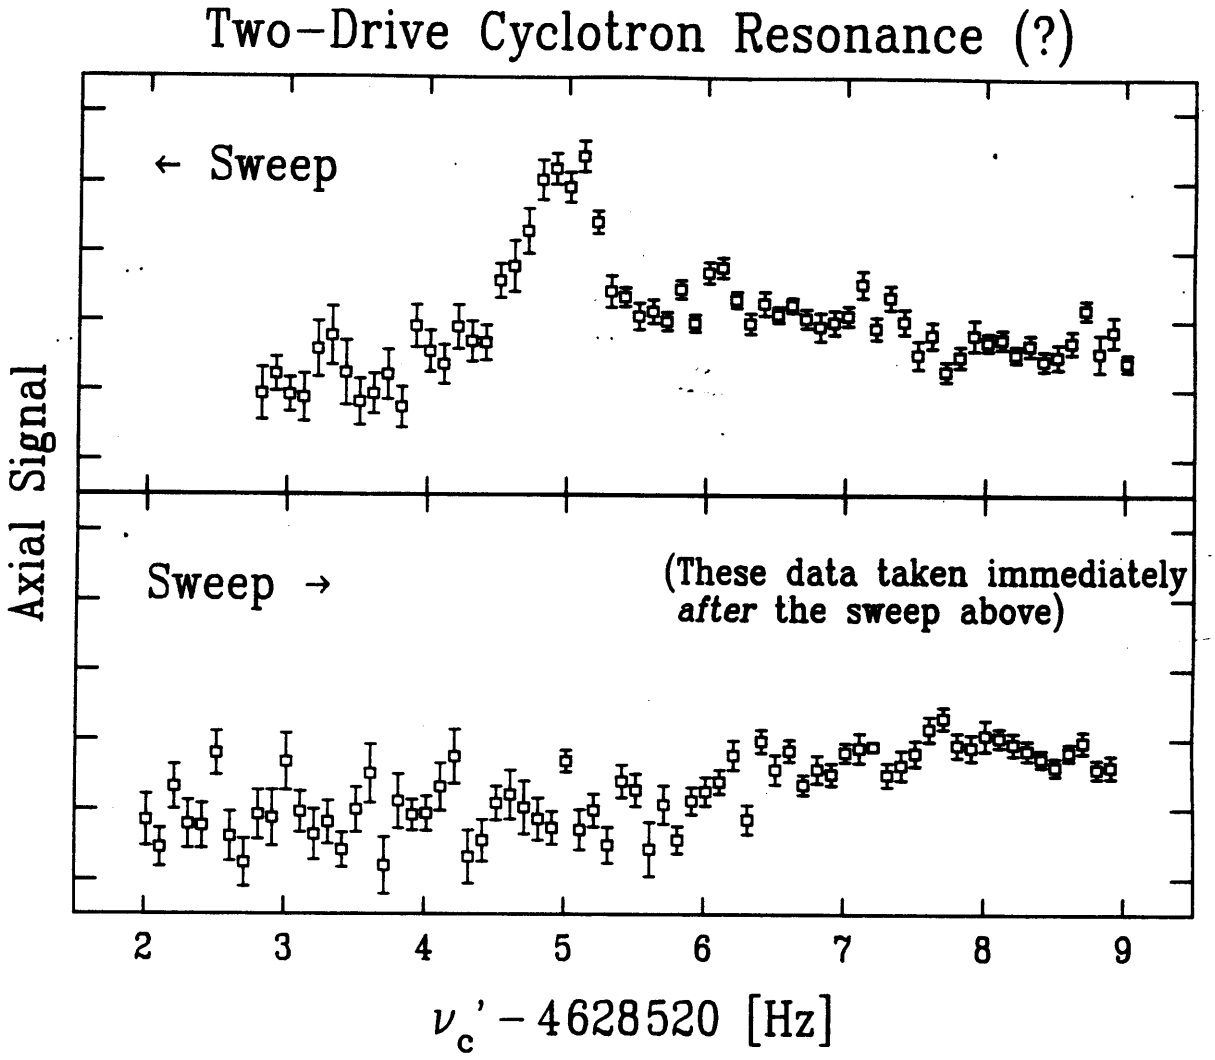

Figure VI.B.3. The subtleties of two-drive cyclotron resonances. The top curve (swept first, right-to-left) shows the resonance and then the increase in non-resonant signal *after* the peak. The lower curve (swept immediately after, left-to-right) shows the increased noise slowly vanishing as we moved back across the cooling resonance.

$\nu_c' - \nu_z$. The frequency difference between the two oscillators was kept fixed at the (previously determined) axial frequency. Each point on Figure VI.B.4 represents a 10-second average of the coherent axial amplitude induced at the *fixed* ν_z. After each average, we moved both oscillators by 0.1 Hz and began averaging anew. The error bars on the graphs indicate the standard deviation of the mean of axial signals averaged. That is, they are not traditional error bars, but rather an indication of how much "noise" (meaning signals not at ν_z) was present at each frequency.

The order and direction of the sweeps is important here. The top curve was taken first, sweeping from right to left. As we reached the middle of the curve, just as the amount of coherent signal appeared to be peaking, the amount of "other" signal increased dramatically. We could not determine whether this extra noise was due to the initial cyclotron motion (left over from the ions' creation) or, possibly, to coherent motion left in the sweep because the cyclotron oscillation was so undamped that we could not sweep slowly enough to assure adiabaticity across the resonance.

However, we have just started induced cyclotron measurements, and at this stage it is too early to tell if these problems can be overcome simply by a more clever choice of experimental parameters. As we shall discuss in the next Chapter, these difficulties have also led us to design other techniques to measure the cyclotron resonances, all variations of separated oscillatory field methods [RAM56], which, in the long run, we think are the techniques most likely to produce the highest precision results.

CHAPTER VII

CONCLUSIONS FOR THE FUTURE

This final chapter summarizes the measurements and calculations we have made to understand systematic effects in the apparatus, especially those parameters that will affect a high-precision mass measurement. The chapter also discusses the directions of our future work.

In the first section, VII.A, we present measurements of some of the perturbations in the trapping fields. We explain several techniques we have used to measure the total electrostatic anharmonicity. Using these methods, we have measured how well the guard rings can compensate for this anharmonicity and, in addition, we measured the (inconvenient) shifts in the natural resonance frequency which accompanied this compensation.

We also explain how we measured the anti-symmetric electrostatic potential used to shift the ions within the trap. Although we have only measured to high accuracy the product of the linear and cubic coefficients of this potential, we use other measurements to estimate (at lower precisions) the individual coefficients. These coefficients also determine the natural damping rate of the ions and the current they induce in the trap.

We also include in Section VII.A evidence for the presence of charged patches on the trap electrodes. Measurements show that after we coated the electrodes with a carbon colloidal paint, these effects were cut down at least by a factor of three.

Section VII.A concludes with brief discussions about trap tilt, eliminating contaminant ions; and magnetic field effects. The discussion of magnetic field problems is brief not because the problems are less important—in fact, they are crucial—but because have only started making careful measurements. The discussion of contaminant ions is brief because our observations are rather qualitative.

In Section VII.B, we outline several modification in the experiment which we are in the process of making, or expect to make in the near future. Some of these improvements—switching to a quieter SQUID detector, eliminating the magnetic bottle—are rather concrete, while others—single sideband detection, separated oscillatory field measurements—deal more with techniques. We will also mention some of the more speculative modifications we are considering: for example, stabilizing the magnetic field and redesigning the trap electrodes to make patch effects less of a problem.

Finally in Section VII.C, we assess the possible accuracy and precision of the experiment, both with the present apparatus and with the modifications we feel will be required in order to reach our goal of accuracy at the 10^{-11} level.

VII.A Measurement of Systematics

In this section, we discuss measurements of various systematic shifts, drifts, and geometric coefficients for ions in the trap. In some cases we measured things which we'd prefer were absent (anharmonicity, patch effects); for other effects, we need to

know the coefficients to interpret other data correctly (anti-symmetric "shifting" potentials); for still others we will need for corrections in our ultimate mass measurements (trap tilt, magnet drift).

We will first discuss the electrostatic potentials. Since most of our experiments focussed on detecting the axial motion, we have made many measurements of electrostatic systematic effects. Using the terminology of Chapter II and Chapter V, we have measured the lowest order anharmonic coefficient, C_4, and the two lowest order anti-symmetric coefficients B_1 and B_3. We have also observed an anti-symmetric potential which persists in the limit of zero applied potential. We attribute this potential to patch effects.

Using the Guard Rings: Measuring C_4

Measurements of the electrostatic potentials, in particular, played a rather symbiotic role with the single-ion experiment. In order to detect one ion, we needed to excite it to rather large orbits in the trap. At these orbits, anharmonicity plays an increasingly important role. On the other hand, these large orbits are the best laboratory for making anharmonicity measurements. Thus, as we were able to see fewer and fewer ions, our ability to characterize the trapping potential improved dramatically, and we therefore could improve the harmonicity to see even fewer ions.

As we discussed in Section II.B, the symmetric electrostatic potential provides axial confinement of the ions. As we showed in Equation (II.B.2), we can write the

symmetric potential as a power series:

$$\Phi = \frac{1}{2} \sum_{k\,\text{even}} C_k \left[\frac{r}{d}\right]^k P_k(\cos\theta) \qquad \text{(VII.A.1)}$$

The $k = 4$ term in this potential is the leading-order anharmonic term. As we showed in Equation (II.B.18), this term leads to an amplitude dependent shift in the axial resonant frequency:

$$\frac{\Delta\omega_z}{\omega_z} = \frac{3}{4} C_4 \frac{z^2}{d^2} \qquad \text{(VII.A.2)}$$

where z is the size of the axial amplitude and d is the characteristic size of the trap, (II.A.2.5). We also showed (Figure II.B.2b) that this shift can lead to bistability; that is, swept resonances can show sensitivity to sweep direction and hysteresis. In Chapter V, we gave some typical examples (Figures V.D.3 and V.D.4) of anharmonic axial sweeps. We will discuss in this section how we used similar data to make quantitative estimates of the anharmonicity; in particular, how we measured C_4.

All the techniques we used were essentially applications of Equation VII.A.2 above. For small numbers of ions (< 10), our best results came from cw swept, two-drive resonances (described in Section V.C); though I will also mention a newer, one-drive pulse method that shows promise of yielding values more quickly. For larger clouds, we had to use more approximate techniques, and I will not mention those methods below.

When the drives are sufficiently large ($V_d \gg V_d^{crit}$ in Equation II.B.21) the

hysteresis can be so large that we should only observe the ion response when we sweep in one direction. For these drives, the anharmonic frequency shift, Equation VII.A.2, determines the total response. We can therefore turn those resonant sweeps "on their side;" that is, plot ω against the measured ion response, z:

$$\omega = \omega_0 + \frac{3}{4}\frac{\omega_0 C_4}{d^2} z^2 \qquad \text{(VII.A.3)}$$

The curvature of the resulting parabola thus yields C_4. The danger of this technique is that it relies on the absolute calibration of the detector and an accurate count of the number of ions in the trap to determine from the SQUID's output the absolute ion response. Errors in those determinations will be squared in the final value of C_4. (However, especially after the single-ion measurement, we have reasonable confidence in our absolute calibrations.)

When the anharmonicity is smaller (or the drives are lowered if we have enough signal-to-noise), we can observe the ions in both direction of the sweep, though hysteresis is still present. (Figure V.D.3 is an example of one such sweep.) For these sweeps, we can use the different peaks in the two directions to determine C_4. When the total frequency shifts involved are of the same order as the (harmonic) width of the resonance, γ_z, we need to be a little more careful when we compute C_4. Using the complete anharmonic line-shape, Equations II.B.19 and II.B.20, one prescription that yields C_4 is:

$$\frac{3}{4}\omega_0\frac{z_0^2}{d^2}C_4 = \frac{\omega_1 - \omega_2 \pm \frac{\gamma_z}{2}\frac{\sqrt{\beta_1^2 - \beta_2^2}}{\beta_2}}{\beta_1^2 - \beta_2^2} \qquad \text{(VII.A.4)}$$

where ω_1 is the frequency at which the maximum response occurs (the *shifted* peak), β_1 is the response at that peak (in units of the trap size, that is, z_{peak}/z_0), and ω_2 and β_2 are the frequency and response respectively at any other point on either of the two sweep directions. For this second point, we usually used the peak on the response curve for the sweep in the "bad" direction: that is, the smaller peak.

Both of these methods (parabola fitting and explicit computation) also allowed us the measure the natural frequency, ω_0, though it is usually very close to the smaller peak. This determination, though, does not depend on the explicit calibrations because both the shifts and our estimate of C_4 depend on the calibration in the same way. Thus these estimates are a bit more robust.

Using these two techniques, we measured the anharmonicity of several clouds of ions at different guard ring settings. Figure VII.A.1 shows the combined results of two of the most precise measurements, separated in time by several months. One measurement (Figure V.D.3 was one of the guard ring settings) was performed on a single N_2^+ ion. The other measurements were performed on a cloud of four N_2^+ ions. Since we changed several things between these measurements, (for example, we entirely disassembled the trap, painted its surfaces with Aquadag, etc.) the absolute guard ring settings were different for the two clouds. To include both clouds on the

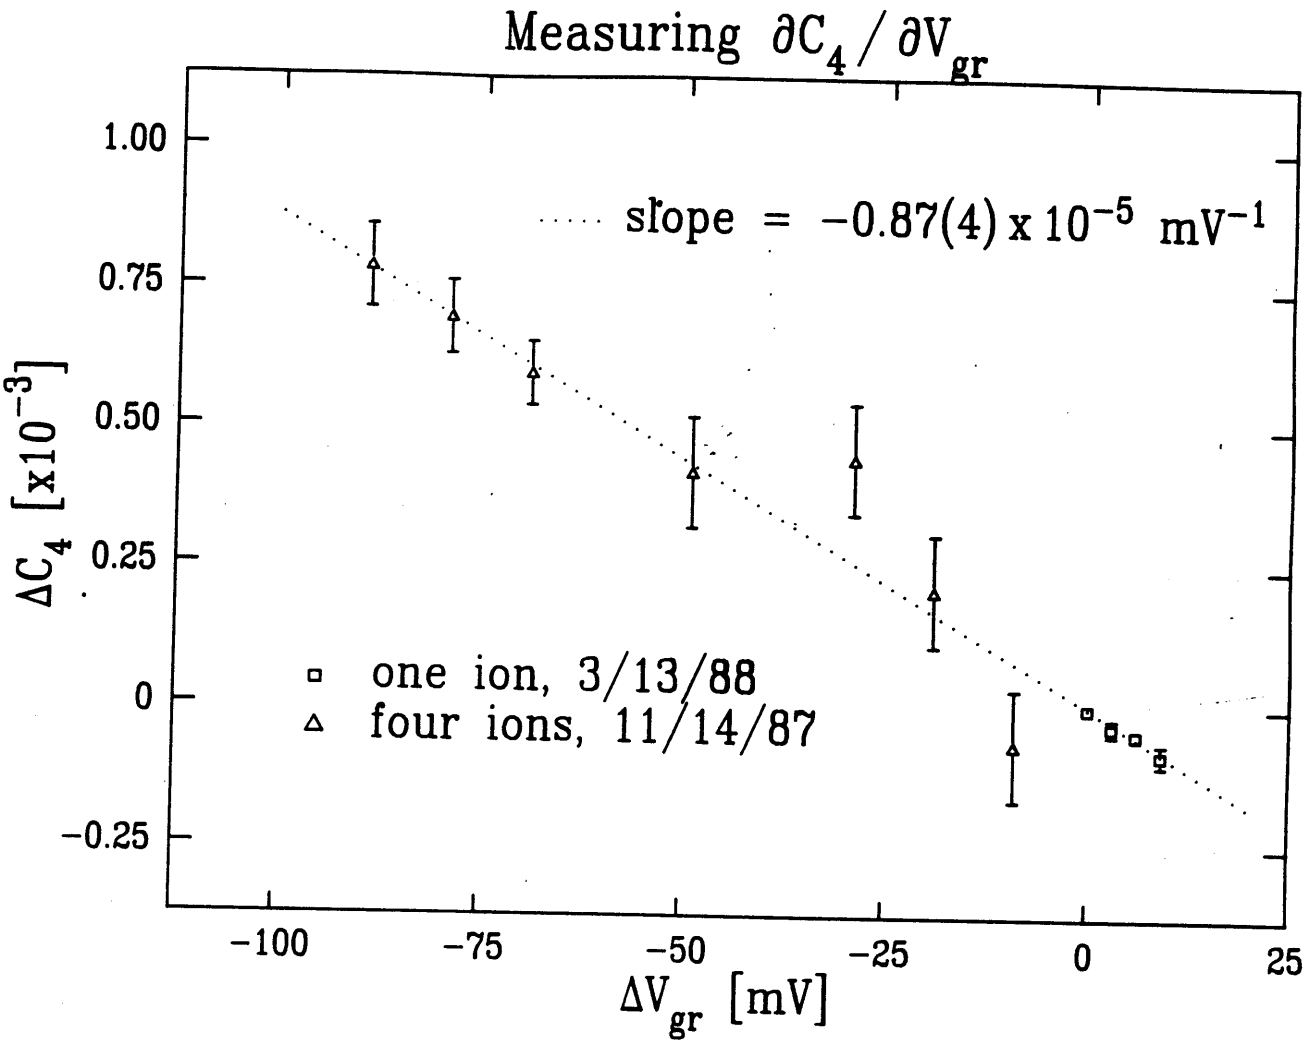

Figure VII.A.1. Guard ring effectiveness. The change in the anharmonicity, C_4, is plotted against the change in the guard ring voltage, V_{gr}.

same graph, then, we have plotted the C_4 measurements against the difference between the guard ring voltage and the voltage at which $C_4 \approx 0$ for each cloud.

As we have discussed in several places, for equal signal-to-noise measurements, the ions show a number dependence of $n^{-3/2}$ in sensitivity to C_4. Thus with one ion, we could make about 8 times better estimates of C_4. Of course, we also needed an 8 times better C_4 just to see the ions at all.

The slope of the points in Figure VII.A.1 gives the change in C_4 for a change in the guard ring voltage. We can express this value as a dimensionless constant by multiplying it by the trap voltage. The resulting constant is the change in C_4 for a given change in (V_{gr} / V_{trap}). This ratio, called D_4 by Gabrielse [GAB83], measures the overall sensitivity of the guard rings:

$$D_4 \equiv V_{trap} \frac{\partial C_4}{\partial V_{gr}} = .076(4) \qquad \text{(VII.A.5)}$$

Gabrielse's paper predicts D_4 for different guard rings shapes at different distances from the center of the trap. For our trap (Figure IV.B.1), his formulae predicts $D_4 = 0.11$. Since D_4 depends exponentially the position of the guard rings (because of the strong electrostatic screening by the nearby ring and endcap), our measured results shows fair agreement with his calculation.

In order for a single ion to show no hysteresis when driven to 20% of the trap size, this value of D_4 requires setting the compensation potential correct to about 1 mV; a difficult, but not impossible demand. However, sweeping these resonances is

rather time consuming, and thus we have been exploring quicker methods. One particular method, using the single-drive pulse scheme of Section V.A shows great promise.

The rapid method to measure C_4 is to measure the axial frequency change of a pulsed ion as a function of its decaying amplitude. To use this method, we pulse the ions to very large orbits, up to half the trap size. After the direct detector excitation decays away, we record the ion's transient decay. We then partition that decay into smaller time chunks and fourier transform these chunks independently. In this way, we measure the ion's frequency at different times, and thus as the ions decays, at different amplitudes. Presently, we can measure the frequencies fairly well, but we do not have enough signal-to-noise to measure the amplitudes with sufficient precision to use them as an accurate indicator of the ion's amplitude. Since anharmonicity shifts the ion's frequency with respect to the detector, these shifts also change the damping time of the ion *during* the transient. Thus we have not been able (yet) to convert the frequency-vs-time graphs to the frequency-vs-amplitude data we need to compute C_4. However, we expect to solve these problems in the near future.

Using the Guard Rings: Shifting ω_z

Whenever we use the guard rings to change C_4, we simultaneously shift the ions' resonant frequency, too. In terms of Equation VII.A.1 above, this means that the guard rings shift C_2 as well as C_4. Although not catastrophic, these shifts are

somewhat inconvenient, and as mentioned in Section II.B, several traps have been designed that should minimize this additional shift. For a simple Penning trap, Gabrielse showed that the ratio of the endcap spacing to the inside ring diameter, z_0/ρ_0, determines the relative shift in C_2 for a given change in C_4. (Different ratios set the ring and endcaps at different distances from the asymptote of the hyperbolic electrodes. Thus the various terms in the guard ring's potential at the center of the trap will be screened differently.) Following his prescription, we constructed our trap with $\rho_0 = 1.16 z_0$ for which there should be no shift in the ions' frequency as we adjust the guard rings.

Using the same raw data that provided the $\partial C_4/\partial V_{gr}$ measurement in Figure VII.A.1, we determined the ions' axial frequency at the different guard ring settings. Again, to combine the two sets of data, we measured the shift in the axial frequency for the different clouds relative to the axial frequency at which $C_4 \approx 0$. The results are shown in Figure VII.A.2. The shifts, though small, are clearly non-zero.

To compare these data with other traps, Gabrielse suggests another dimensionless parameter, D_2, that measures the shift in the ions' frequency that the guard ring produces relative to the shift that an equal voltage change on the ring would have produced. For the data in Figure VII.A.2, the trapping potential was 8.725 V and the resonant frequency was about 158.9 KHz. Using Equation (II.A.2.5), then, a 1 mV change on the ring would produce a 9.1 Hz shift. Thus using the slope of Figure VII.A.2:

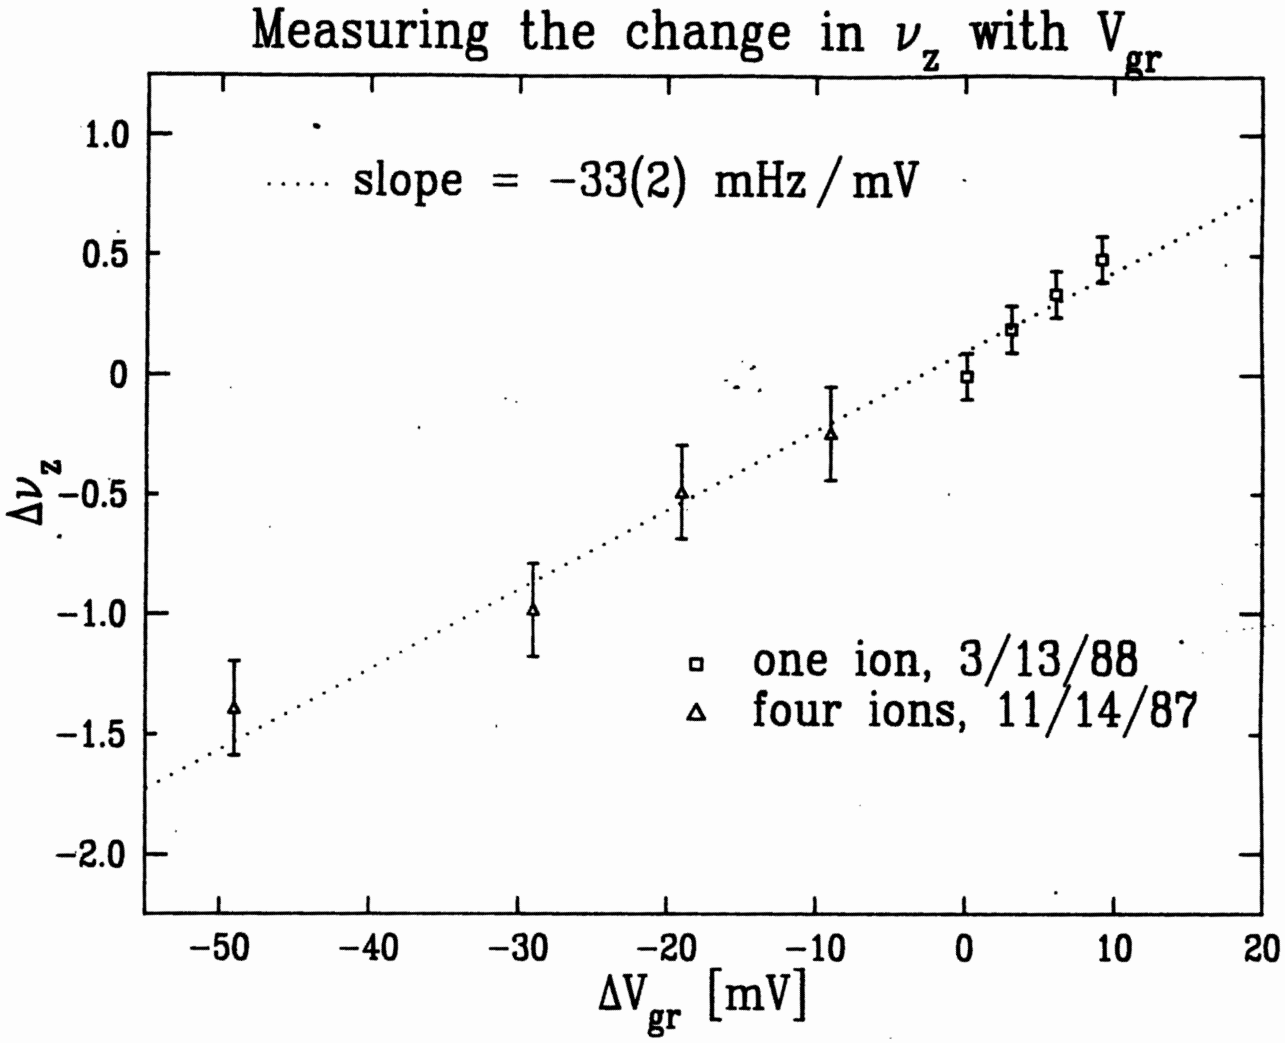

Figure VII.A.2. Guard ring optimality. The change in the resonant frequency caused by a change in the guard ring potential is plotted against that guard ring change.

$$D_2 \equiv \frac{\partial\omega_z / \partial V_{gr}}{\partial\omega_z / \partial V_{trap}} = 3.7(2)\times 10^{-3} \qquad \text{(VII.A.6)}$$

As a measure of trap-tuning convenience, Gabrielse suggests an "optimality" coefficient, γ, which is the ratio of the two dimensionless D-constants just discussed. For example, moving the guard rings further away would decrease D_2, but such a decrease would be meaningless because the guard rings' ability to cancel out anharmonicity would be diminished equally. This ratio γ thus measures (approximately) the relative change in C_2 for a given change in C_4. The smaller this ratio, the more conveniently we can tune out trap anharmonicities.

For our trap, comparing Equations VII.A.5 and VII.A.6 we see that:

$$\gamma \equiv \frac{D_2}{D_4} = 0.049(4) \qquad \text{(VII.A.7)}$$

This value should be compared to the γ typical for non-optimal traps. The usual practice (before GAB83 and BEA86) was to build "asymptotically symmetric" traps; that is, traps for which the ring and the endcaps were symmetric about their mutual asymptote: $\rho_0 = \sqrt{2}z_0$. All these traps, regardless of the location of the guard rings had $\gamma = 0.56$. Thus our trap shows better than a factor of 10 improvement in optimality, similar to other traps in current use. [VMF85]

We have one rather disturbing piece of historical data. We had made crude measurement of the shifts caused by the guard rings with large clouds of ions even before we could measure C_4 accurately. All of these measurements show a three times larger D_2. It is very difficult to find a mechanism that can explain these results. For

example, we could get C_4 shifts in the axial frequency due to the magnetron orbit size. (Equation II.C.4.10) For a constant magnetron orbit size, then, changing C_4 will also change ω_z. However, using the measured sensitivity D_4, a magnetron orbit of 1 mm would be required to cause such a large shift. Though not impossible, such a consistently large magnetron orbits seems improbable. We have no other reasonable explanation for this discrepancy.

Before concluding our discussion of anharmonicity, it is worth recalling that C_4 is not the only anharmonic term. For example, asymmetric charged-patches or machining errors (ruining reflection symmetry) could cause odd order potentials. In addition, we expect some higher-order anharmonicity due to the symmetric terms, C_6, C_8, etc. In both cases, then, we need to extend (VII.A.2) to include terms that look like z^4/d^4, and so on. Therefore, even though we may be able to eliminate entirely C_4 we none-the-less might have amplitude-dependent shifts. Alternatively, and a bit exotically, if we have large patches, we may have adjusted C_4 to cancel a substantial C_3 term in the z^2/d^2 term (which also contributed at that order like C_3^2 [LAL76]) only to left with a rather large z^4/d^4 term in which these contributions add together.

Regardless of their origin, we certainly have observed the effects of these higher order terms. For example, Figure VII.A.3 shows a very anharmonic single-ion sweep (swept right-to-left) taken at the same time as the single ion data for the previous two figures. The C_4 term (that is to say, the z^2/d^2 term in the frequency shift) was

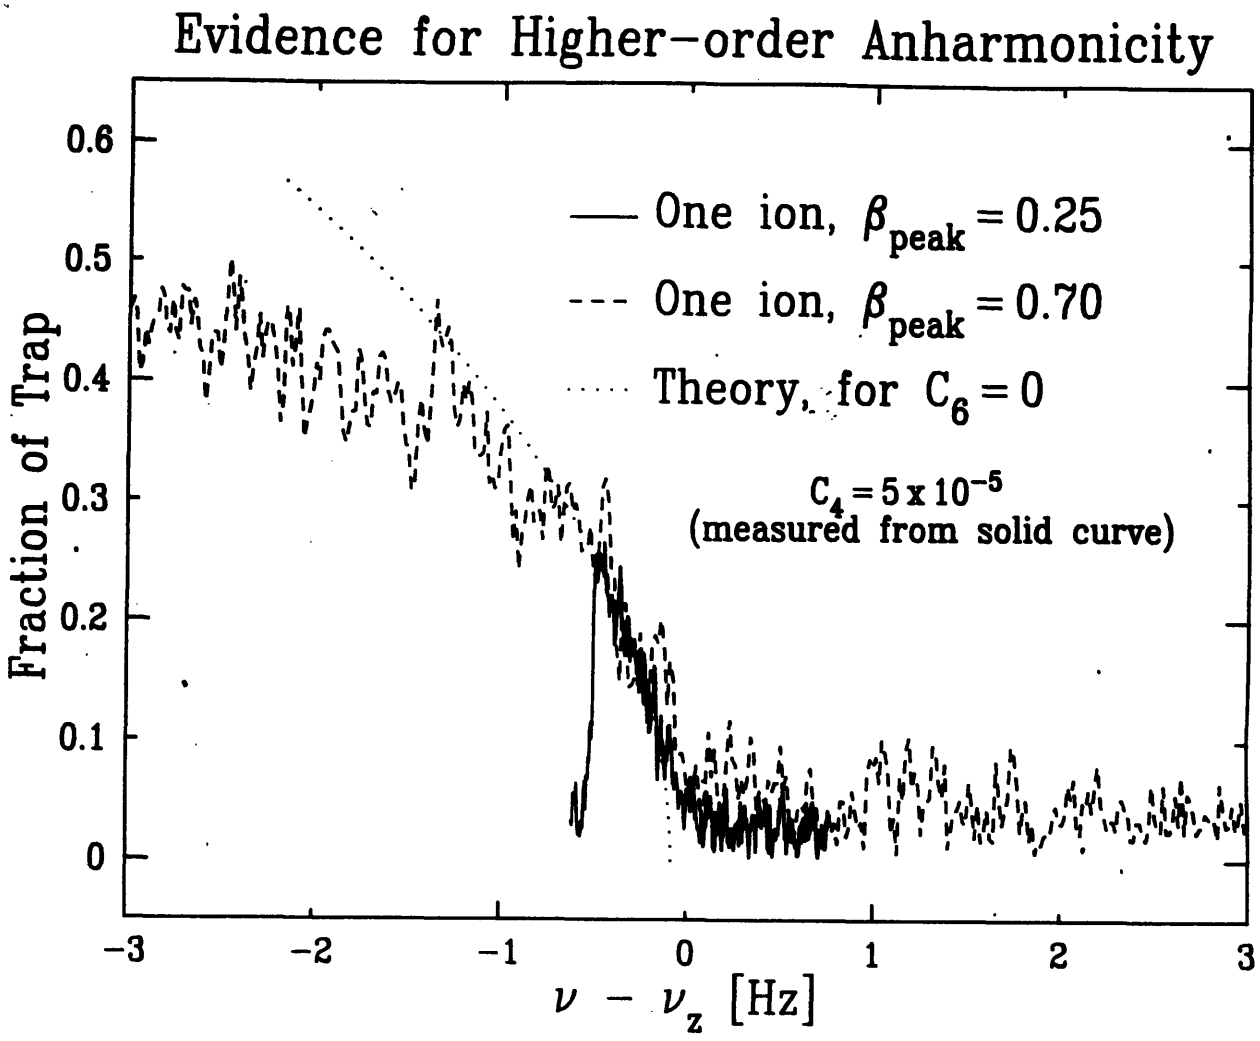

Figure VII.A.3. Demonstration of higher-order anharmonic effects in the trap. The solid line shows an anharmonic sweep of one N_2^+ driven to 25% of the trap size. The dotted line extrapolates the measured anharmonicity for stronger drives. The dashed line, taken at the stronger drive, shows that C_4 alone is insufficient to account for all the anharmonicity the very strongly driven ion experienced.

measured using an excitation which drove the ion to about 25% of the trap size. (the solid line of Figure VII.A.3) We then increased the drives by a factor of 3 (the dashed line). When the ion's motion was smaller than 30% of the trap size, the ions response was identical. However, as the sweep continued, we clearly see that the ion's resonant frequency is shifting faster than the square of its amplitude. (The dotted line on Figure VII.A.3 is an extrapolation assuming the ion only been subject to a z^2 shift.) Thus using very strong excitations, we have strong evidence for other orders of anharmonicity.

Measuring B_1 and B_3

In Chapter II, we described in detail a method to shift the ions in the trap by applying voltages to the lower endcap. By investigating the anti-symmetric potential in the trap:

$$\Phi = \frac{1}{2} \sum_{k\,\mathrm{odd}} B_k \left[\frac{r}{z_0}\right]^k P_k(\cos\theta) \qquad \text{(VII.A.8)}$$

we discovered that the cubic term caused shifts in the resonant frequency of the ions as we shifted a cloud within the trap:

$$\nu_{eff} = \nu_z - \frac{3}{4}\frac{d^4}{z_0^4} B_1 B_3 \left[\frac{V_-}{V_r}\right]^2 \qquad \text{(VII.A.9)}$$

where V_- is the voltage on the lower endcap and V_r is the ring (trapping) voltage. (V_+ has been set to zero.)

Thus we can measure the product $B_1 B_3$ by measuring the resonant frequency as

we put different voltages on the lower endcap.

Figure VII.A.4 present the data from one such measurement. To obtain these data, we used the one-drive pulsed method of Section V.A on a cloud of about 20 N_2^+ ions. We used this method because it provides the quickest measurement of the ion's frequency at the precision required for these data. After each measurement, we shifted the voltage on the lower endcap and, to keep the ions at the same resonant frequency, we shifted the ring voltage by half that amount. (Recall from Section II.B that changing just the lower endcap voltage changes both the symmetric and anti-symmetric potentials.) We also shifted the guard rings so that they were at the same fraction of the ring voltage. We used a 5 1/2 digit DVM to set the trap potentials. Since the trap voltage was about 9 Volts, we could therefore only make reproducible changes in the resonant frequency to a bit better than 1 Hz. (This coarseness in the frequency allowed us to use the "sloppier" one-drive pulse method to determine the frequencies.)

We fit these data to a parabola, shown as a solid line on Figure VII.A.4. Using Equation VII.A.9 above, to convert the curvature of this parabola to a value for B_1B_3, we found $B_1B_3 = 0.198(4)$ where the error bar is the 1σ limit from the non-linear χ^2 fit. We have made five different measurements of B_1B_3 in the past nine months. Except for one measurement, (made right after we cooled the trap, and when there may have been a background gas problem), all the measurement agree fairly well with this results. Averaging the four remaining measurements yields:

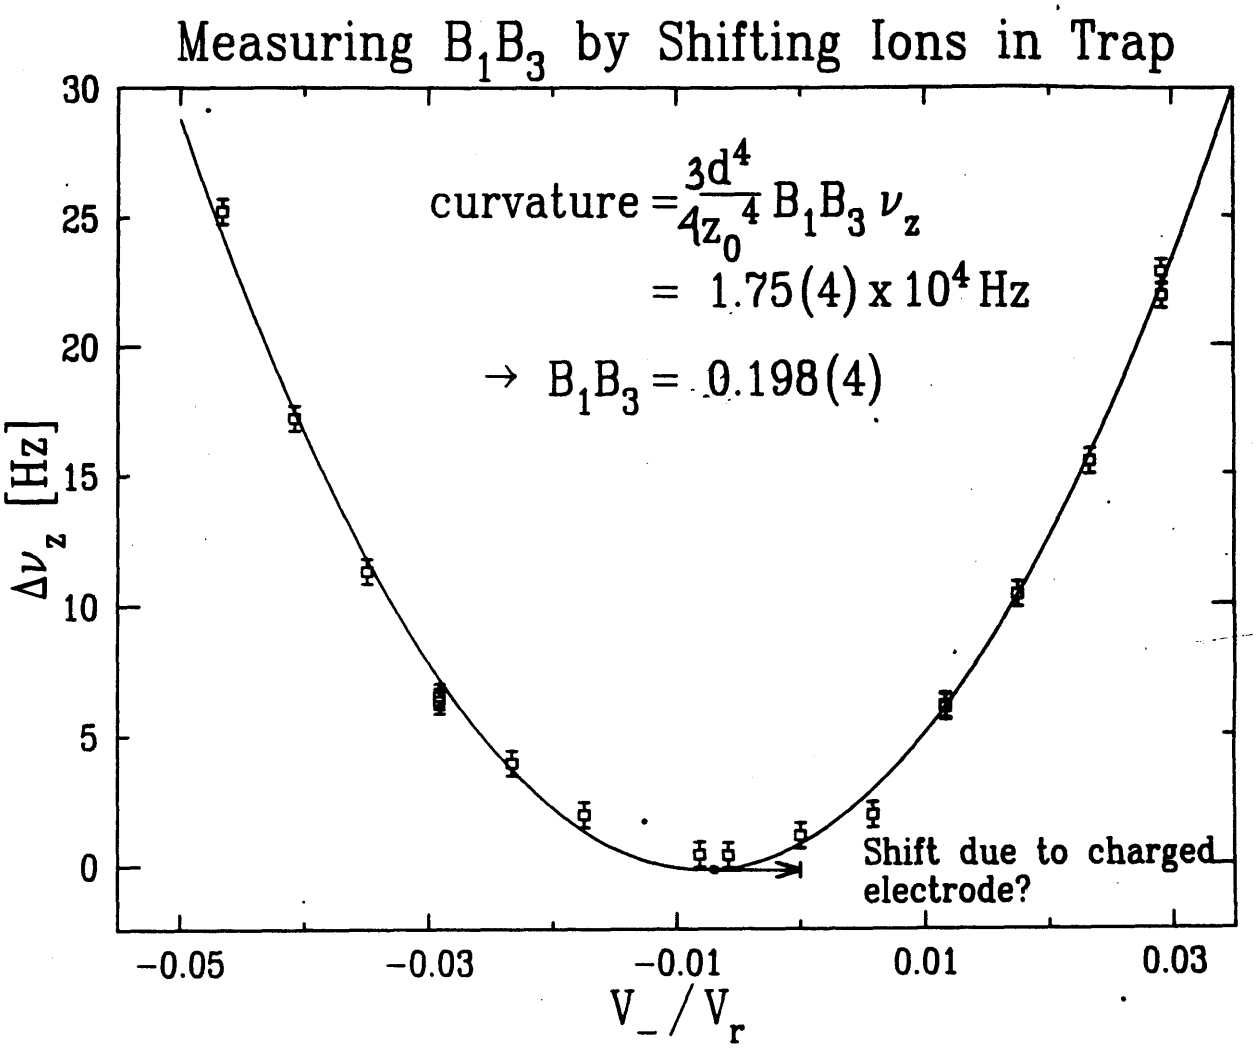

Figure VII.A.4. Measuring the anti-symmetric potential in the trap. The equilibrium position of the ions were shifted by a potential applied to the lower endcap. The frequency shift occurs because of a cubic term in the electrostatic shifting potential.

$$B_1 B_3 = 0.202(5) \qquad \text{(VII.A.10)}$$

This value agrees quite well with an estimate from a graph in the Brown and Gabrielse review article [BRG86], which predicts this value to be about 0.21.

We noticed in all our early measurements that the parabola was always displaced, indicating the presense of an additional, anti-symmetric voltage on one of the endcaps. For example, for the data in Figure VII.A.4, the shift was 68(3) mV. There are many possible causes for small differential potentials, including thermocouple effects in the lower endcap bias wire, and different work functions for the different endcaps. From the measurement above on a single ion species, we couldn't rule out the possibility that machining errors had caused a shift in the center of the trap. The 68 mV error above corresponds to a 5 mil shift in position, whereas the axial dimension was specified to better than 1 mil [FLA87], so trap asymmetry seems unlikely to be the major cause. Quite possibly, the major part of the asymmetric potential comes from built-up charge on the surfaces; that is, to a patch effect. We present further evidence that points to a patch effect in the next sub-section.

Although we could measure the product $B_1 B_3$ with fairly high precision, it was quite difficult to isolate either of these coefficients. We discussed one crude way in Section II.B. The endcap voltage at which the ions are shifted out of the trap, V_-^{kill}, depends on the sum, $B_1 + 3B_3$, as described in Equation II.B.13:

$$B_1 + 3B_3 = \frac{z_0^2}{d^2}\left[2\frac{V_{trap}}{V_-^{kill}} - 1\right] \qquad \text{(VII.A.11)}$$

(This expression takes into account the fact that V_- both shifts the ions and weakens the trapping potential.) We measured V_-^{kill} several times. For example, to make the single-ion plateaus (Figure V.D.1), we edged V_- closer and closer to V_-^{kill} in order to release one ion from the trap at a time. However, we never observed a "hard" limit; that is, some ions would leave the trap at V_- well above (0.2 V) other ions. (A possible explanation is given below.) Thus we could only set a fairly rough limit on V_-^{kill}. For the trap voltage set at 8.72 V, the most we can say is that $V_-^{kill} > 7.36$ V. When this value is used in (VII.A.11) above, we find that $B_1 + 3B_3 < 1.64$. This limit can be combined with our value for B_1B_3 above. The resulting quadratic equation sets the (rather relaxed) limit:

$$\begin{aligned} 0.19 < B_3 < 0.36 \\ 1.06 > B_1 > 0.56 \end{aligned} \qquad \text{(VII.A.12)}$$

though physically, B_1 cannot be greater than one.

We have other, indirect measurements of B_1. The amplitude of the single ion response in Figure V.D.1), for example, depends on B_1 as discussed in Section II.A, Equation II.A.3.4. Using our detector calibration and our drive calibrations, the single ion peak in Figure V.D.1 sets $B_1 = 0.80(5)(8)$, where the first error bar is statistical from the detector noise, and the second is (a guess, really) of the calibration reliability. A somewhat more sensitive determination of B_1 comes from the measurements of the

few-ion damping constants in Figure V.B.1. We can use the three values of γ_z measured from those data, to make a better estimate of γ_z^0, the damping of a single ion on the detector's resonance: $\gamma_z^0 = 0.31(2)\, s^{-1}$. Using the room temperature measurement for the detector inductance, and the frequency and Q of the detector at liquid helium temperatures, we can invert (II.A.3.7) to give B_1:

$$B_1 = \frac{2z_0}{e}\sqrt{\frac{m\gamma_z^0}{\mathrm{Re}\,Z}} \qquad \text{(VII.A.13)}$$

$$= 0.77(2)(4)$$

where the statistical error comes from γ_z^0 above, and the other error represents a (conservative) 10% uncertainty in the cryogenic inductance. This final value of B_1, combined with the measurement of $B_1 B_3$, gives $B_3 = 0.26$. These final estimates agree well with the [BRG86] calculation.

However, if these estimates are correct, then an interesting effect may be happening. We see from Equation (II.B.11) that for large V_-, ω_{eff} can become zero: our trap disappears, and our ions will leak away. We can use (II.B.11) and (II.B.8) to determine the value for $\hat{z}$ (the new equilibrium position) at which we lose trapping. If this value is less negative than $-z_0$, then the ions will "de-trap" before they are pushed into the lower endcap. We find:

$$\hat{z}^{de-trap} = -z_0\left[\frac{1}{6B_3}\frac{z_0^2}{d^2}\left\{2\frac{V_{trap}}{V_-^{de-trap}} - 1\right\}\right] \qquad \text{(VII.A.14)}$$

Comparing this expression to Equation VII.A.11 above, we see that the trapping potential vanishes just as the ions reach the lower endcap for $\frac{B_1}{B_3} = 3$. For larger ratios, the ions will hit the lower endcap; for smaller ratio, the trapping potential will vanish before they get that far.[1] We see that our best estimates for B_1 and B_3 are very close to this limit. This effect may explain the relative "softness" of the V_- barrier mentioned above. Our trap may be close enough to the limit ratio that the thermal motion of the ions is enough to get them over the vastly weakened potential barrier when we adiabatically relax[2] the trap with increasing V_-.

Patch Effect

As mentioned above, the B_1B_3 measurement demonstrated an offset in the anti-symmetric potential. To better understand these offsets, we put different ion species in the trap and measured the potential required to bring these ions into resonance with our detector. Inverting Equation II.A.2.5, we see that:

$$V_{trap} = \left\{\frac{d^2\omega_0^2}{e}\right\}m \qquad \text{(VII.A.15)}$$

Thus the mass of the ions and the trapping potential should be directly proportional.

1. However, since most of their velocity is in the z-direction, the ultimate fate of the ions will still likely be smashing into the lower endcap.
2. The speed of the relaxation is essentially unimportant: when we lower the potential, we shift the ions off the detector, and they become uncoupled from the thermal bath.

Even in the presence of machining errors (see above), we expect a graph of V_{trap} vs. m should go through the origin.

Figure VII.A.5 shows two such measurements of trapping voltage versus ion mass, made with two different detectors, and thus at two different frequencies. In both sets of data, the three species lie precisely on a line, and this line does *not* pass through the origin. In both cases, the curves show about a 150 mV offset at zero mass, though differing by about 10 mV. (The total offset is many times greater than even the most conservative error estimate.) The magnitude of this offset precludes thermocouple effects, though not some of the more exotic, low-temperature metal-metal offsets.

We can use the slopes of the lines in Figure VII.A.5 as a measurement of the trap size, d^2. Using the known frequencies and the known charge and mass of the ions, we found, for the dashed line:

$$d^2_{146250} = 0.2988(2)\ \text{cm}^2 \qquad \text{(VII.A.16)}$$

and for the dot-dashed line:

$$d^2_{159600} = 0.3001(2)\ \text{cm}^2 \qquad \text{(VII.A.17)}$$

where the subscripts on d^2 correspond to the frequency of the detector. The difference in length between the two values of d is less than 0.5 mil. Since the trap was cycled to room temperature between these measurements, we find this a very satisfactory agreement. The theoretical size of the trap, from the machining drawings, should have

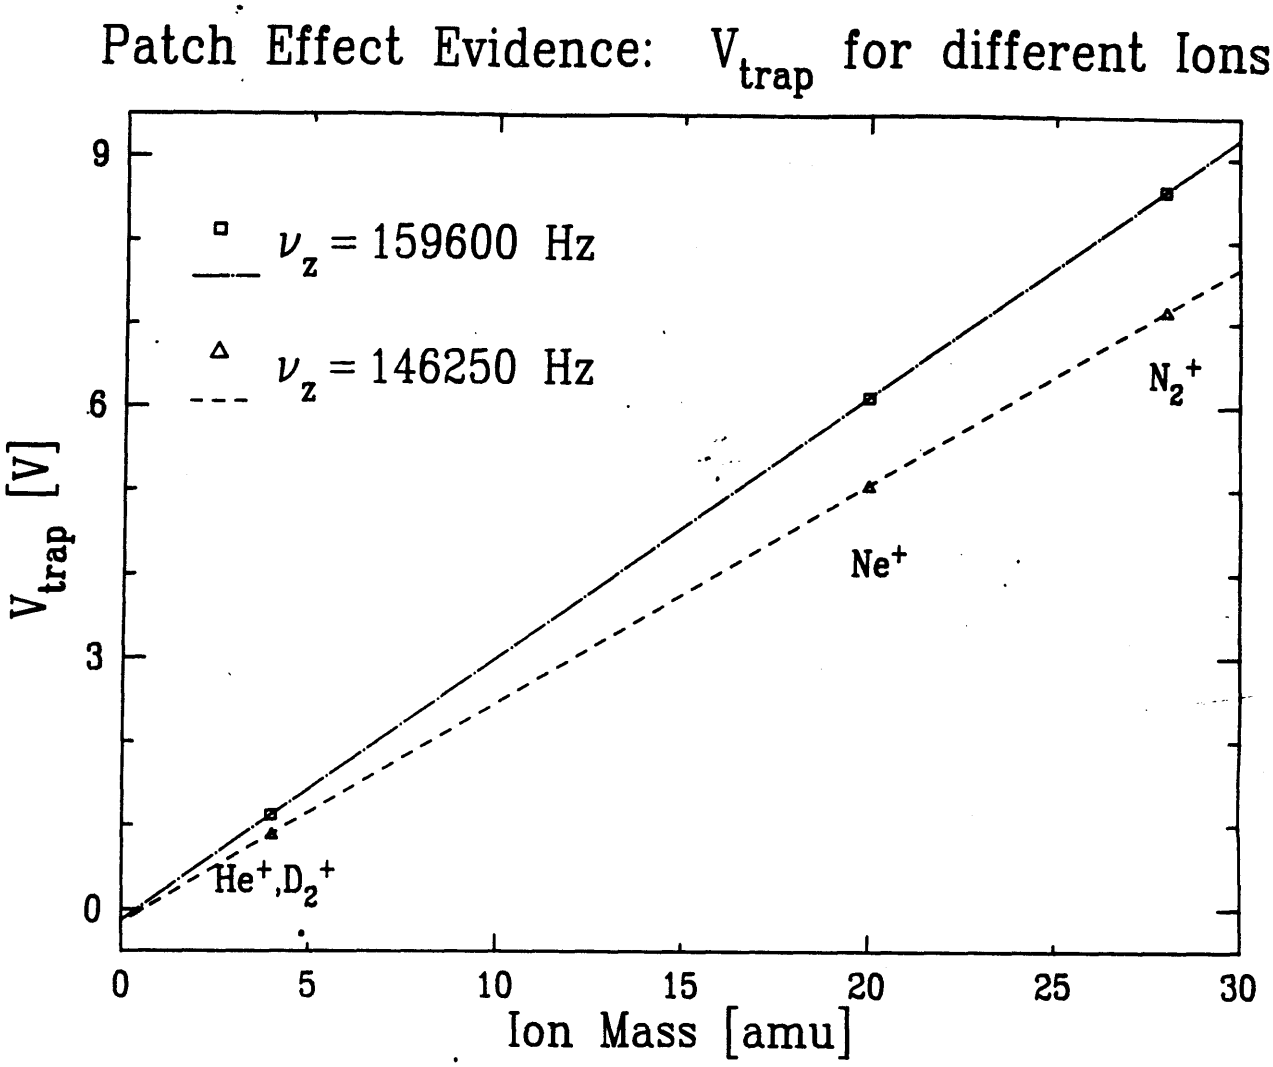

Figure VII.A.5. Trapping potential required to bring different species of ions into resonance with our detector. The two, independent measurement have y-intercepts of about −150 mV, indicating a differential voltage on the trap surfaces.

been $d^2 = 0.3011$ cm^2; again agreeing with our measured value to about 0.5 mil.

After we coated the trap with carbon paint, we performed another B_1B_3 measurement. The results, shown in Figure VII.A.6, show a factor of three improvement in the offset of the parabola: the solid line in Figure VII.A.6, is offset is 19(3) mV from the origin. Also included on this Figure are two earlier measurements, taken several months apart during which the apparatus was kept at liquid helium temperature. (One measurements was presented as Figure VII.A.4.) We have included these data to show that the voltage offsets had been very stable. Since the voltage offset was reduced by coating the surface of the trap, we must conclude that the offset was caused by a charged patch.

We have not yet performed a test of voltage vs. mass like that of Figure VII.A.5. The reason is practical: to make a measurement at low masses, we need use a gas, Helium or Deuterium, neither of which pump very quickly from the trap region. Therefore, we usually make these measurements right before we warm the apparatus to make other modifications. However, combining our one data point (N_2^+ at 8.725 V) with the measured values of d^2 above, we find about a 50 mV offset, again a factor of three improvement over the unpainted trap. While this improvement has helped in the detection of a single, N_2^+ possible, we may find that we need additional improvements to go to lower masses, such as $^3He^+$.

Trap Tilt

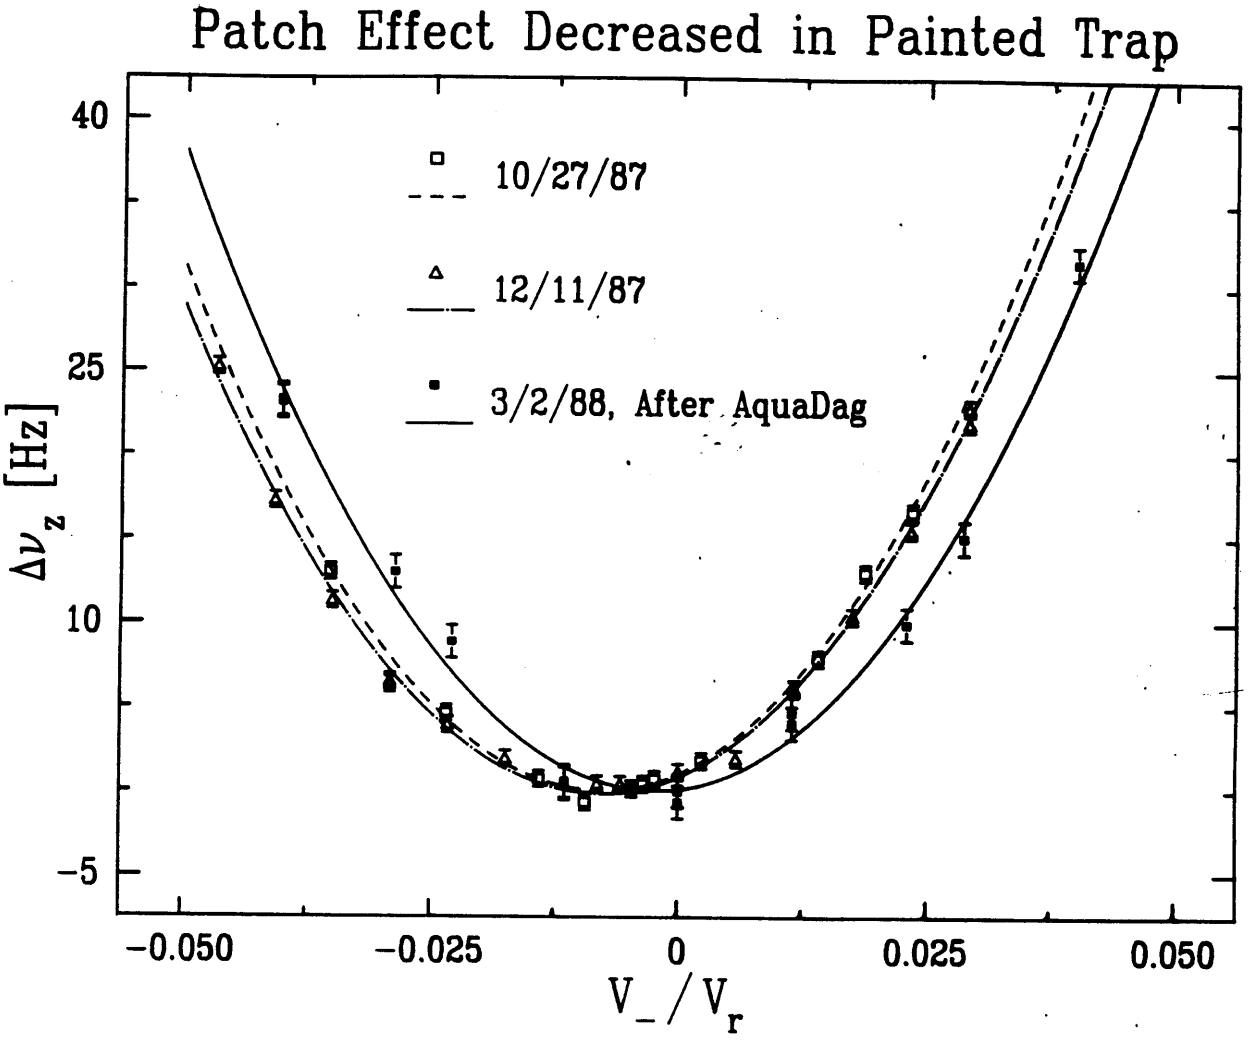

Figure VII.A.6. After the trap was coated with carbon paint, we observed a decrease in the voltage offset, here apparent as a shift in the location of the parabola minimum.

We discussed in Section II.C.7 the effects of misaligning the magnetic and electrostatic axes of the trap. The signature of this tilt is that the usual relationship, $\bar{\omega}_z^2 = 2\bar{\omega}_c\bar{\omega}_m$, will not be satisfied. (The bar denotes measured values of the ions in the trap) Equation II.C.7.7 gives a prescription to convert the trap frequencies into a tilt. In the small angle limit, we find the tilt, θ:

$$\theta \simeq \frac{2}{3\bar{\omega}_m}\sqrt{\bar{\omega}_m - \tilde{\omega}_m} \qquad \text{(VII.A.18)}$$

where, as in Section II.C.7, $\tilde{\omega}_m \equiv \frac{\bar{\omega}_z^2}{2\bar{\omega}_c}$.

We used this method to measured the tilt before we last warmed the apparatus. At that time, $\bar{\nu}_z$ = 159594 Hz, $\bar{\nu}_c$ = 4628076 Hz, and $\bar{\nu}_m$ = 2756(1) Hz. (The magnetron frequency was determined from magnetron cooling data, for example, Figure VI.A.2) Using these data, we see that $\tilde{\nu}_m$ = 2751.7 Hz, and thus, using (VII.A.18), we find:

$$\theta = 1.5(3)\ ^\circ \qquad \text{(VII.A.19)}$$

When we warmed the experiment (primarily to apply Aquadag to the trap), we physically measured the tilt using a level and some shim stock. We found a 1–2° tilt, in agreement with the measurement above. When we reassembled the trap, we did our best to zero this tilt. We believe that the tilt now should be less than 1°. Although measurements of the cyclotron frequency have not been made, a more accurate measurement of the magnetron frequency, combined with an extrapolation for the

cyclotron frequency from the magnet drift (see below) indicate a tilt of about 0.8°. As discussed in Section II.C.7 (see Equation II.C.7.9), this tilt would cause a negligible ($< 2\times10^{-14}$) error if we use only the measured cyclotron and axial frequencies to compute the true cyclotron frequency of $^3He^+$, ignoring the magnetron frequency measurement entirely.

Killing the Bad Ions

Qualitatively, one of the worst nuisances continues to be contamination by other ions. Impurity ions could cause shifts, anharmonic effects, time-dependent drifts; all appearing as irreproducibilities in the measurements. When there enough impurity ions in the trap, we have had difficulty detecting the good ions at all! For example, when we had a high background pressure of Helium gas in the trap, we frequently made so many helium ions that we could not detect any other species. (This problem was traced to the ion pump, which had been saturated with Helium from an earlier cryogenic leak. We removed this pump from the system.)

However, even in the absence of leaks, we still expect to create several undesired ion species at the same time that make the ions we wish to trap. The Penning trap is so deep and mass insensitive that any background neutral ionized within its confines will remain trapped. We have tried many techniques to eliminate these ions, and though no one of them seems to work by itself, in combination they seem to remove the contaminant ions for many hours, at least from the central region

of the trap.[3]

The primary difficulty in eliminating the contaminant ions is that at large orbits, the anharmonicities in the trap shift the ion frequencies unpredictably, especially right near the surface of the trap. Thus simple resonant techniques are incapable of adding enough energy to expell all the bad ions. However, driving very strongly near the (calculated) resonant frequency for each contaminant species[4] followed immediately by dropping the lower endcap voltage to drive all the ions near (but not *too* near) the lower endcap has proven fairly successful. The driven ions should have large orbits and thus should hit the lower endcap at the lower excursion of their excited orbit. The undriven ions, ideally unexcited, should have orbits too small to bring them into contact with the lower endcap.

It was hard for us to believe that such a seemingly fool-proof method should fail, but it does. Typically, three or four hours after the treatment above, the axial resonance begins shifting and we need to repeat the process. Perhaps ion/ion collisions convert some of the axial motion of the contaminant ions into magnetron motion while they are being driven, and thus give these ions large radial orbits. At

3. As mentioned in Section II.C.2, we could tell that the good ions were being affected by contaminant ions when we saw their resonant frequency begin to drift. The magnitude of these drifts were typically of $\sim Hz$ / min order.
4. We usually assumed water, $^{14}N^{+}$, tungsten (from the field emitter), gold (from the surface), helium, and hydrogen would be the most likely contaminants.

large radii, they are unaffected by the drive and cool very slowly, presumably by coupling to the good ions in the center of the trap.

More recently we have tried a variation of the method above, using *incoherent* excitations. We use the computer to generate white noise (to an adjustable bandwidth) and use this voltage to amplitude modulate an excitation drive set near the contaminant ions' resonant frequency. By clever choice of parameters, we can put the good ions at a node of the noise (it has a sinc^2 dependence far away from its central frequency) to minimize their excitation. Again, we dip the cloud after such an excitation. Using this technique, we have trapped a single good ion, apparently with no other contaminant ions. (For example, one ion was stable for over 13 hours, when we let it go free.) However, to achieve this isolation, we sacrificed many of the good ions each time we shifted the cloud, indicating that the good ions were still being heated during the contaminant ions' excitation.

Magnetic Field Measurements

As we shall discuss in Section VII.C below, magnetic field perturbations—drift, inhomogeneity, and short-term instability—will be the most important factors that determine the accuracy of our mass measurements. However, at this point, we have made only minimal observations of these perturbations. In fact, we only have data for one of these factors—the magnetic field drift—and those data should be considered preliminary.

In Figure VII.A.7, we show our four reliable cyclotron measurements. The frequencies were obtained from the measured cyclotron and axial frequencies, corrected according to Equation II.C.7.9. The error bars are statistical, and, importantly, do not contain any systematic corrections for orbit sizes. Since the corrections for the magnetic bottle alone are likely to be at least as large as the error bars given, (see Sections II.C.5 and VI.D.) the spread of these points about the least-square line is not a surprise. For these data, $\chi^2 = 8$, showing that we have likely underestimated the errors at least by a factor of two.

We have made no measurements of the magnetic bottle, or verified the radial cooling limits in Section VI.A. Our earlier calculation [FLA87] for the bottle caused by the copper and MACOR in the trap recently was repeated, confirming the earlier results: the trap should cause about a 6 gauss/cm^2 magnetic bottle.

To measure this bottle, we could move the ion cloud within the trap (using the lower endcap) and measure the cyclotron frequencies at the different locations. Since we have fairly good values for B_1 and B_3, we should have a fairly good idea of the absolute equilibrium position of the shifted cloud. Alternatively, we could measure the cyclotron frequency with the axial motion locked at different amplitudes (using the absolute calibrations of the detector to determine the absolute amplitude of the axial motion) and measure the shifts in the cyclotron frequency, (II.C.5.5) to measure the magnetic bottle. This last method may be particularly relevant because it measures directly the systematic corrections we will need to apply in a real cyclotron

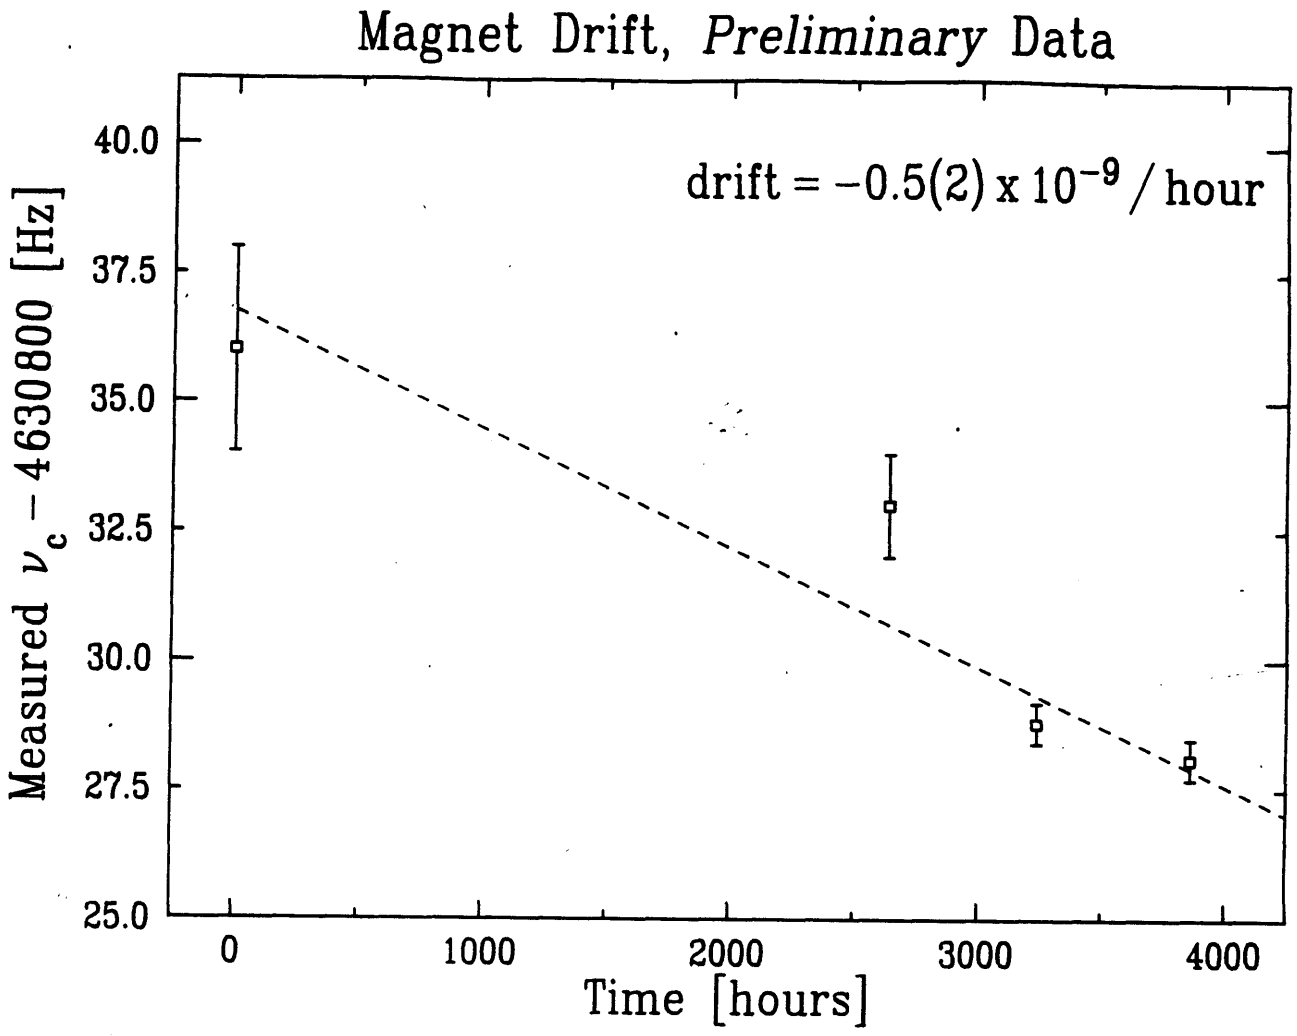

Figure VII.A.7. Preliminary magnet drift measurements. The cyclotron frequencies are plotted against time. These data have not been corrected for bottle effects.

measurement.

To verify the magnetron cooling limit, anharmonicity effects will probably be easier to use than the magnetic bottle. Although the magnetic bottle also shifts the magnetron frequency (sometimes called the "continuous Stern-Gerlach" effect [BRG86]), these shifts can be overwhelmed by electrostatic anharmonicity. Comparing their shifts (Equations (II.C.4.4) and (II.C.5.4)), using the magnetic inhomogeneity calculated above, an electrostatic anharmonicity of only $C_4 = 4\times10^{-6}$ would cause an identical shift to our present bottle.

Thus one method to measure the magnetron radius would be to use our guard ring calibrations to change C_4 and measure the axial frequency shifts beyond those given in Figure VII.A.2. However, these shifts probably will be too small to measure a radius anywhere near the magnetron cooling limit. (For example, a single-ion line width, 0.05 Hz, would be caused by a radius $\rho_m = 0.25$ mm at $C_4 = 10^{-4}$, about forty times larger than the expected cooling limit, (VI.A.16).

Alternatively, then, we could lock the axial frequency (as discussed in Section V.D) with C_4 purposely set rather high—but not so high that the resonance is bistable—and then measure changes in the locking voltage when we heat the magnetron motion from its cooled state. As the magnetron orbit size increases exponentially, the trap voltage will have to shift to accommodate the changing axial resonance frequency. Using the observed growth rate of the magnetron orbit, we

should be able to extrapolate back to determine the original (cooled) magnetron orbit. This procedure would also let us verify the cooling and heating rates given in Section VI.B.

VII.B Proposed Improvements

In this brief section, we shall outline some of the modifications we believe will improve the accuracy and precision of the experiment. As we shall see in the next section, several of these improvements will be required in order to reach our desired precision of 10^{-11}. Others mentioned here will simply make the goal more convenient.

One obvious area of improvement is the SQUID detector. Following our discussion in Section III.A.2 (the J. Appl. Phys. paper), there are three possible directions for improvement. If we can increase the Q of the detector, we will lower its thermal noise on resonance, and thus improve our ion current sensitivity. We also have a higher-frequency SQUID (which uses 200-MHz RF pump frequency) This device shows noticeably quieter voltage noise (at least, when it is not connected to the ion trap), and thus, all other things being equal, we could increase the coupling between our coil and the SQUID to increase the bandwidth of our detector.

In addition, we have been using pulsed excitation and transient detection more and more frequently. However, we have optimized the detector for CW use, minimizing its total, effective current noise on resonance, i_n. However, for pulsed experiments, the integration time is limited by the ion's life-time. Since the life-time

depends inversely on the real part of the detector impedance, we therefore must minimize $i_n^2 \mathrm{Re}\, Z$, the effective noise *power*. Likely, then, we will find that we want to operate off the detector resonance (where $\mathrm{Re}\, Z$ has decreased but i_n^2 is still relatively small), and tighten the coupling to recover the lost signal.

Another area of improvement will be in the electric and magnetic field stability. Although I leave a more complete discussion until the next section, it is worth emphasizing our unique requirements for short-term magnetic field stability: instability in the magnetic field translates directly into inaccuracy in measured mass ratios. Electric field instability and anharmonicity, though less crucial, are far from unimportant. For example, although we have reduced the patch effect, we have not eliminated it. We even may need to sputter a new surface onto the trap which is less sensitive to patch formation than gold. Alternatively, we have discussed redesigning the trap in rather radical ways. For example, we might construct a trap with the endcaps very close together while the ring remains further away. The trapping frequency depends on the combination of these dimensions ($\omega_z \sim d^{-1}$, $d^2 = z_0^2/2 + \rho_0^2/4$), so such a trap would not have vastly different frequency properties. On the other hand, patches on the rings would have minimal effects, whereas the endcap electrodes would have much smaller surface area to support patches. However, this kind of trap will require significant theoretical work (for example, to determine how to compensate it efficiently, etc.) and thus we have no short-term plans to implement these ideas.

We have several schemes to improve or eliminate the magnetic bottle. One idea, for which the design is complete, to decrease the magnetic bottle by about a factor of five by adding MACOR disks at strategic locations. We could use the superconducting shims of our magnet to eliminate the remaining inhomogeneity. However, the extra MACOR would significantly lengthen the thermal time constant of the trap, and with all the magnetic compensation disks, it might take a week (or longer) for the trap to cool to 4 K. Another option would be to use different materials as insulators in the trap (for example, fused quartz) which has better thermal, conductivity and magnetic properties.

Finally, we have been developing better detection techniques and technology. One simple improvement will be to use single-sideband detection techniques. Presently, for both transient and phase-sensitive detection, we use a double-balanced mixer to bring the ion signal into a convenient frequency range for the computer to measure and store. However, after mixing with $\cos\omega_{mix}t$, both a signal at $\omega = \omega_{sig}$ and noise at $\omega = 2\omega_{mix} - \omega_{sig}$ will appear at a detected frequency $\omega_{sig} - \omega_{mix}$. That is, by mixing with only one phase, $\cos\omega_{mix}t$, we cannot distinguish the sign of the resulting frequency. For example, signals above and below ω_{mix} will show up in the FFT spectrum.

However, by using two mixers, driven independently by $\cos\omega_{mix}t$ and $\sin\omega_{mix}t$ we can recover the complete input signal. For example, we can use the two mixed signals as the real and imaginary parts of a complex time signal and perform the

complex fourier transform on the resulting data set. Exploiting the fourier-transform symmetry relationship, $\tilde{X}(\omega) = \tilde{X}^*(-\omega)$, we can reconstruct the positive and negative spectra from the even and odd parts of the resulting complex spectrum. Off the detector resonance, where the noise is due entirely to the voltage noise of the SQUID, this technique should provide a $1/\sqrt{2}$ improvement in the amplitude signal-to-noise, providing somewhat less improvement on the detector's resonance.

A second improvement in techniques is to switch to pulsed, separated field measurements, akin to those devised by Ramsey [RAM56,PRI86]. In addition to yielding a $\sqrt{2}$ improvement in peak-center determination (by narrowing the resonance line), separated field measurements are particularly useful on undamped systems for which sweeping across the resonance with CW excitation is impractical.

These methods all use the ion's free precession between pulses to give an amplitude response that is phase dependent. As an example, we describe one experiment which has already yielded results. We first excited the magnetron mode with a frequency-selective pulse near the magnetron frequency. After waiting some time, T, we pulsed the ion with a magnetron-axial pulse. The product of the amplitude and the duration of that pulse was designed to swap the action between the axial and magnetron motions. (See Section VI.A) Thus the phase of the resulting axial motion (which we detected) should depend sinusoidally on the free-precession period, T. Using this method, we have measured the magnetron magnetron frequency to 10 mHz; a factor of 100 improvement over our CW cooling method. (This was

how we determined the magnetron frequency for the recent tilt measurement described above.) We expect to attempt this method on the cyclotron resonance in the near future.

VII.C Future Prospects and Future Limits

Finally, having demonstrated that we could detect one trapped ion, having shown several techniques for measuring cyclotron frequencies which soon will be attempted on single ions; having measured or estimated many of the systematic perturbations to the cyclotron frequency, we shall conclude in this section by speculating on how well we might be able to compare masses in our experiment, both with our current apparatus and ultimately.

There are two classes of problems that limit the accuracy of this (or any high-precision) experiment: run-to-run instabilities, and measurement imprecision. We have two major kinds of measurement uncertainties; first, errors when we extrapolate to zero perturbation and second, our ability to split thermally-broadened, noise-contaminated resonance lines. In this section, we will discuss first the limits causes by magnetic and electric field instabilities. Many of these limits can be improved by actively or passively stabilizing the fields. We then turn to measurement imprecision. Using estimates of the perturbations, thermal motions and detector signal-to-noise, we try to give predictions for the best width of cyclotron lines, our ability to split those lines, and our ability to combine several measurements to extrapolate to zero perturbations. We will find that the thermal motion and the zero-drive extrapolations

are intimately related.

Magnetic Field Fluctuations

The first source of field fluctuation comes from the slow decay of the magnetic field. All superconducting magnets show some drift [WIL83], and even though our drift is particularly small (measured to be 10^{-9} per hour or so), we could imagine that errors in measuring this drift will translate into uncertainty in the field at the 10^{-11} level, if other resports using similar magnets (VMF85) prove consistent.

Much of the additional field fluctuations could come from changes in the ambient conditions. When the atmospheric pressure changes, for example, we expect several small changes in the apparatus. First, the pressure change should directly alter the size of the magnet. An extremely crude calculation, assuming that the magnet has the same bulk modulus as steel gives, for typical pressure fluctuations in the Boston area, instabilities somewhat below level of 10^{-9} per hour level. More importantly, though, [GAB88,CAS88] changes in atmospheric pressure change the boiling temperature of the liquid helium surrounding the magnet, and thus can cause thermal expansion or contraction of both the magnet and our ion trap, again at the 10^{-9} level. To overcome these problems, like other experimenters [GAB88], we probably will need to regulate the pressure of our Helium bath against an absolute pressure. This regulation, though crude, has performed well enough to bring these pressure-induced variations below measurable levels; that is, below 10^{-11}, and, we expect, could be

extended even further.

External magnetic field "noise," caused by the small fluctuations in the Earth's field and by magnetic pollution (due to trucks, trains, elevators, and low frequency electromagnetic fields which propagate though ferromagnetic building materials) may prove to be an even worse problem. Although superconducting magnets provide (via flux quantization) some inherent screening against external field changes [GAT88], we probably will have to perfect some additional field stabilization methods. At Washington, Van Dyck's group [VMF86] have begun using a "variable bottle," (two cleverly chosen loops of superconducting wire) which, among other things, screen out external magnetic field changes by keeping the flux threading the loops constant. This method screens out environmental fluctuations as well as the field fluctuations of the superconducing magnet. Another method, currently being implemented at Harvard by Gabrielse [GAT88], is to add extra superconducting solenoids along the axis of the magnet which should increase by a factor of ten or more the passive self-shielding of the magnet system without sacrificing the field homogeneity. To keep fields stable to the 10^{-12} level, we probably will need to combine these techniques, possibly with some form of active stabilization.

On the other hand, if we can *simultaneously* trap the two ions whose masses we wish to compare (an open theoretical question), many of these field problems could be reduced considerably, and, in fact, we might get to the 10^{-11} level using any of the existing stabilization methods.

Electric Field fluctuations

Electric field fluctuations are important, too, because the axial frequency affects the measured cyclotron frequency, ω_c' through the magnetron shift, Equation (II.A.2.13):

$$\omega_c' \approx \omega_c - \frac{1}{2}\frac{\omega_z^2}{\omega_c} + \cdots \qquad \text{(VII.C.1)}$$

Therefore fluctuation in the axial frequency $\delta\omega_z$ become fluctuations in the measured cyclotron frequency $\delta\omega_c'$:

$$\delta\omega_c' \approx -\frac{\omega_z}{\omega_c}\delta\omega_z \qquad \text{(VII.C.2)}$$

When we measure all the modes simultaneously, these fluctuation are not a problem: the quadrature sum rule (II.C.7.4) takes care of them. (Shifts in the magnetron frequency are rarely a problem because there must also be a large tilt for them to have an impact on the cyclotron frequencysee further the discussion in Section II.C.7, especially equation II.C.7.9.) However, sequential measurements would be a problem. In order to understand the problems electric field instability might introduce, let us describe its effects for a particular measurement scheme.

We expect that "separated-field" measurements have the greatest chance to provide extremely high precision frequency resolution. For such a measurement, we would excite the cyclotron motion directly, wait some extended time, then pulse the coupling drive to transfer the cyclotron motion into the axial motion, swapping action

between the two modes as described in Section VI.A. By measuring the phase of the resulting axial motion as a function of the time between the excitation and coupling pulse we could reconstruct the natural cyclotron frequency. The precision of this kind of measurement depends roughly on the product of the signal-to-noise and the maximum length of time we can allow the ion to freely precess. The field instabilities therefore limit the precision by limiting the time we can integrate the phase during the free precession.

We can estimate the electric field stability using preliminary measurements we have made on pulsed, single ions. Though only a rough estimate, we saw fairly reproducible stability in ω_z to about 1/20 Hz over a four second period, consistent with estimates for the noise of the voltage reference given in Chapter IV. (Pressure-induced temperature fluctuations changing the size of the trap should be two orders of magnitude smaller [FLA87]). Assuming the axial frequency performs a kind of random walk, we might naively expect a roughly $T^{3/2}$ dependence for the phase error in time, T. (One factor of $T^{1/2}$ for the random walk, one factor of T for the integration over frequency to get the phase.) Using this instability estimate, we would be limited to a 40 second integration time for N_2^+ and about 160 seconds for $^3He^+$, thus yielding a minimum width $\frac{\Delta\omega_c}{\omega_c} \approx 10^{-9}$ for N_2^+ and 2×10^{-11} for $^3He^+$. (We get the longer times for $3He^+$ because its larger ω_c suppresses the axial contribution to a greater extent than in N_2^+.) The somewhat disappointing figure for N_2^+, however, can be overcome

by locking the axial motion to an external frequency source (and possibly locking the ion's amplitude, too) to avoid this diffusion of the axial frequency.

Thus both electric and magnetic field instability will be a problem at the 10^{-9} and active and/or passive stabilization of the fields will be required to push beyond that level. However, the stabilization required should not be particularly sophisticated or complex, and thus these field fluctuations, though extremely important, should not present a *fundamental* or insurmountable limit in the future.

Thermal Widths and Shifts

Each of the major perturbations alters the mode frequencies through the random, thermal motion of the ions. As we saw in Section II.C, the perturbations shift the mode frequencies by the square of the various mode amplitudes. The axial mode is coupled directly to a heat bath through the detector dissipation. Since the cyclotron mode, in turn, must be coupled to the axial mode to measure it, all modes will be given both widths and shifts from this thermal interaction.

In general, we need to separate and understand widths and shifts of four different types in cyclotron motion for each of the major perturbations. First, the axial frequency itself may have shifts, which, as described in Equations VII.C.1 and VII.C.2, indirectly affect the cyclotron frequency. Second and third, the thermal energy in both the cyclotron and axial modes will broaden and shift the cyclotron frequency directly. Finally, the thermal motion in the modes makes it impossible to know precisely the

initial conditions before excitation. When we then extrapolate to zero drive strength, we will have uncertainties because we could not know the true excitation.

Thus, if we plot measured cyclotron frequency versus amplitude of the drives with with we measured it, we expect a parabola because all the shifts are second order in the amplitudes of the modes. The first three effects above yield errors in the cyclotron frequency while the last effect causes an error in the amplitude. The local slope of the parabola at that amplitude then convert this amplitude error into a frequency error.

How large are these effects for the various perturbations? We will show that for the present MACOR trap, the magnetic inhomogeneities should far outweigh the two other dominant perturbations, electric anharmonicity and relativistic shifts.

Electrostatic Anharmonicity

We first show that the electrostatic anharmonicity (Section II.B.3) will have nearly negligible thermal effects. As shown in Equation (II.C.4), anharmonicity presents a fairly minor effect in the cyclotron motion since it is suppressed by the large factor ω_m / ω_c. Even for the heavier ions (like nitrogen), this shift will be several orders of magnitude smaller than the bottle shift. The axial width, given by Equation (II.C.4), could be significant when the trap is not well compensated. Estimating z^2 from equipartition for an oscillator in equilibrium at temperature T, the axial width is be given by:

$$\frac{\Delta\omega_z}{\omega_z} = \frac{3}{4}\frac{C_4}{d^2}\frac{kT}{eV} \approx 4\times10^{-3}\frac{C_4}{m\ (\text{amu})} \qquad \text{(VII.C.3)}$$

where V is the trapping potential and m is the mass of the ion in atomic mass units. Using (VII.C.1), we see that this axial width translates into a cyclotron width:

$$\frac{\Delta\omega_c}{\omega_c} \approx 6\times10^{-9}\, C_4\, m\ (\text{amu}) \qquad \text{(VII.C.4)}$$

and thus is never a problem for the kinds of compensation we have achieved, $C_4 \leq 10^{-5}$. The magnetron shifts, which could be larger on an relative scale, will be suppressed by an even larger factor, and, as discussed in Section II.C.7 only enter the cyclotron equation through a tilt in the trap. We therefore neglect them.)

Magnetic Inhomogeneity

The bottle, on the other hand, can cause a much larger width. First, we see that the thermal motion in the axial mode will cause a width in the cyclotron frequency. Using (II.C.5.5), we see that:

$$\frac{\Delta\omega_c}{\omega_c} \approx 2\times10^{-9}\frac{B_2\,(\text{Gauss}/\text{cm}^2)}{m\ (\text{amu})} \qquad \text{(VII.C.5)}$$

In addition, as discussed above, we expect a breadth in the line because of the uncertain initial conditions. Using a cyclotron/axial coupling pulse that completely swaps the action between the two modes, the cyclotron radius, ρ_c, when transferred, becomes a larger axial peak amplitude, z_{peak}:

$$z_{peak} = \sqrt{\frac{\omega_c}{\omega_z}}\rho_c \qquad \text{(VII.C.6)}$$

(See Section VI.A) We then detect the phase of the induced axial motion. However, while the cyclotron motion is precessing, the random component of ρ_c is causing a small shift. Our uncertainty in the cyclotron orbit size, σ_ρ, propagates into an uncertainty in the shift, σ_{ω_c}:

$$\sigma_{\omega_c} \approx \omega_c \frac{B_2}{B}\rho_c\,\sigma_\rho \approx \omega_z \frac{B_2}{B} z\,\sigma_z \qquad \text{(VII.C.7)}$$

where σ_z is the thermal axial motion.

These two widths (the direct thermal width and the initial-condition-uncertainty width) can have different relative magnitudes for different ions and different size drives. Using the present signal-to-noise of our detector (in terms of an axial displacement, about 0.2 $mm/Hz^{1/2}$) and our best estimate for the bottle (about 6 $Gauss/cm^2$), we can tabulate the bottle shift, the spread in the shift due to the uncertain initial conditions, and the uncertainty in a measurement with a 40 second separation time on a single N_2^+ ion, for several different cyclotron pulses and their resulting axial amplitudes (the cyclotron frequency is about 4.630 MHz):

Axial Amplitude (mm)	Bottle Shift (mHz)	Thermal Spread (mHz)	S/N-Limit Uncertainty (mHz)
1.2	80	8	0.5
0.6	20	4	1.0
0.3	5	2	2

We see that for the high signal-to-noise pulses, the uncertainty in the frequency shift

will be due almost entirely to the uncertain initial conditions, and thus there is little reason to use long separation times on the larger amplitude pulses. We thus see that the bottle limits the possible accuracy of the experiment to around 10^{-10} for N_2^+.

For $^3He^+$, on the other hand, (VII.C.5) predicts a much larger relative width in the cyclotron frequency due to the thermal z motion, about 4×10^{-9}. Thus integration times much longer than one second would be superfluous with our existing trap. In fact, for the Helium/Tritium measurement, the uncertainty due to random initial conditions should be swamped by the width of the bottle. Computing the same table for the $^3He^+$ cyclotron at 43.2 MHz, we get:

Axial Amplitude (mm)	Bottle Shift (mHz)	Thermal Spread (mHz)	S/N-Limit Uncertainty (mHz)
1.2	80	20	20
0.6	20	10	35
0.3	5	5	70

Thus, assuming that we have sufficiently small (and well measured) anharmonicity in the trap, we can use the large pulses effectively to measure the bottle and thus have some hope of performing a measurement at a precision of about 4×10^{-10} (about 1 eV) on the Helium and Tritium.

In addition to the thermal spread, there will be a small thermal shift due to the fact that we neglected the average thermal energy in the cyclotron motion; much smaller than the shift due to axial energy because the thermal cyclotron motion is so much smaller. Even for our large bottle, this effect gives a cyclotron shift of about

10^{-12} (its the same for all ions). In addition, it is only through the *uncertainty* in the temperature of the cloud that this width gives an error, and thus we expect this error ultimately should be less than 10^{-13}.

Special Relativity

However, we have plans to eliminate this magnetic inhomogeneity. Our magnet, before we put the trap in, had nearly 3000 times better homogeneity. [OXF85,FLA87] If our MACOR "shimming" (described in the previous section) works and we can reduce this bottle by a factor of five or so we may be able to use the superconducting shims on the magnet [OXF85] to eliminate the rest of the inhomogeneity. Alternatively, as mentioned in the previous section, we might make a trap from less paramagnetic materials and thus stay within the shimming limits of our magnet without adding extra magnetic material. If we can eliminate the bottle in either of these ways, would we then be home free?

No. As discussed at the end of Section II.C, special relativity takes over when the bottle disappears. Comparing (II.C.5.5) and (II.C.8.1), we see that for an ion of mass m, the shifts in the cyclotron frequency due to the cyclotron orbit size, special relativity behaves just like a bottle of size B_2^{rel}:

$$B_2^{rel} \approx \frac{60}{m^2} \frac{\text{gauss}/\text{cm}^2}{\text{amu}^2} \qquad \text{(VII.C.8)}$$

Conveniently, relativity for $^3He^+$ corresponds to roughly the same size bottle as the one used in the previous table, about 6 gauss/cm^2.

Although we may neglect the thermal-axial width (because relativistic shifts due to axial motion are reduced by the large factor ω_c^2/ω_z^2) we cannot eliminate the uncertainty in the cyclotron shift due to the uncertainty in the cyclotron's initial conditions. This remaining thermal uncertainty is about 10^{-11} for $^3He^+$, and thus for large signal-to-noise pulses, we should thus have a corresponding error less than 1 mHz.

Extrapolating to zero relativistic perturbation, then, we get an uncertainty, roughly, of about 10^{-11} divided by the square-root of the number of times we can measure the ions before the field instabilities overwhelm the uncertainty.

Special relativity coupled with field instability thus provides the fundamental limit to our measurements. However, the specific values above still might be overcome. Recall from Section III.C that the axial and cyclotron temperatures are related by:

$$T_c = \frac{\omega_c}{\omega_z} T_z \qquad \text{(VII.C.9)}$$

and thus the cyclotron temperature is usually much higher than the ambient 4 K. The thermal uncertainty in the cyclotron orbit (given by the cyclotron temperature) could be lowered significantly if we could somehow beat this temperature down.

We could lower the axial temperature, perhaps by using continuous feedback or, since we can measure the axial motion more accurately than its thermal noise, by a clever application of a "super-cooling" pulse with the ion off the detector resonance.

By combining this temperature lowering with axial/cyclotron coupling, we could reduce the uncertainty in the cyclotron orbit size.

Alternatively, we could cool the cyclotron motion directly. For example, we could split the ring, use a variable-frequency tuned circuit to damp the cyclotron mode, and change its tuning to remove the damping during a cyclotron measurement. After cooling the axial motion with *its* detector, we would move the ion off the detector so that it would have a very long equilibration time with the axial detector. If we detune far enough so that the equilibration time is longer than the time we take to measure the cyclotron frequency, the uncertainty in the cyclotron initial condition would then come from a much cooler source. Conceivably, we also could cool the cyclotron motion to 4 K directly by collisions with background ions (which are in equilibrium with a detector), and then dump out all but the good ion before measuring its mass. If we want to be *really* far-fetched, we might imagine using collisions with trapped positrons which would cool quickly to 4 K by their microwave cyclotron radiation! Since the final precision, in general, is directly proportional to the cyclotron temperature, any method that could lower this temperature should yield dramatic improvement.

Summarizing, we must stabilize the magnetic field to achieve precisions beyond 10^{-9}. In addition, our present magnetic field inhomogeneities will limit the experiment at the 10^{-10} level. If we can control these problems, we should be able to reach our goal of 10^{-11} and, if we can cool the motions further, beyond.

REFERENCES

ABS70 M. Abramowitz, I. Stegun, Handbook of Mathematical Functions, New York: Dover Press, 1970.

ARN78 V. I. Arnold, Mathematical Methods of Classical Mechanics, New York: Springer-Verlag, 1978.

BAR78 S. Barbanera, et. al., J Appl Phys, 49, 905 (1978).

BEA86 E. C. Beaty, Phys. Rev., A33, 3645, (1986).

BRG82 L. Brown and G. Gabrielse, Phys. Rev., A25, 2423, (1982)

BRG86 L. Brown and G. Gabrielse, Rev Mod Phys, 58, 233 (1986).

BRI74 E. O. Brigham, The Fast Fourier Transform, Englewood, N.J.: Prentice-Hall, 1974.

CAS88 M. Cassidy (Oxford Instruments), Private communication.

CDL77 C. Cohen-Tannoudji, B. Diu, F. Laloe, Quantum Mechanics, New York: Wiley, 1977.

COR77 A. Corney, Atomic and Laser Spectroscopy, Oxford: Clarendon Press, 1977.

DEH67 H. G. Dehmelt, Advances in Atomic and Molecular Physics, 3, 53 (1967).

DSP79 Digital Signal Processing Committee, Programs for Digital Signal Processing, New York: IEEE Press, 1979.

EHN77 G. J. Enholm, J Low Temp Phys, 29, 1 (1977).

FLA87 R. W. Flanagan, Jr., Ph.D. Thesis, M.I.T., 1987.

FLS64 R. P. Feynman, R. B. Leighton, M. L. Sands, Lectures in Physics, Volume 3, Reading, MA: Addison-Wesley, 1964.

GAB83 G. Gabrielse, Phys. Rev., A27, 2277, (1983).

GAB84 G. Gabrielse, Phys. Rev., A29, 462, (1984).

GAB87 G. Gabrielse, Private communication.

GAB88 G. Gabrielse, Private communication.

GAD85 G. Gabrielse, H. Dehmelt, Phys. Rev. Lett., 55, 67, (1985).

GAT88 G. Gabrielse, J. Tan, J. Appl. Phys., in press.

GDK85 G. Gabrielse, H. Dehmelt, W. Kells, Phys. Rev. Lett., 54, 537, (1985).

GIH78 R. P. Giffard, J. N. Hollenhorst, Appl Phys Lett, 32, 767 (1978).

GOT79 K. Gottfried, Quantum Mechanics, Vol. 1, Reading, MA: Benjamin/Cummings, 1979.

GUR77 A. V. Gusev and V. N. Rudenko, Sov Phys JETP, 45, 637 (1977).

HAR78 F. J. Harris, Proc IEEE, 66, 51 (1978).

HOG80 J. N. Hollenhorst, R. P Giffard, J Appl Phys, 51, 1719 (1980).

HOH80 P. Horowitz, W. Hill, The Art of Electronics, New York: Cambridge University Press, 1980.

HPE77 J. Hough, J.R. Pugh, W. A. Edelstein, W. Martin, J Phys E, 10, 993 (1977).

HST49 Hipple, Sommer, and Thomas, Phys. Rev., 76, 1877, (1949).

HUL21 A. W. Hull, Phys. Rev., 18, 31, (1921).

JAB75 L. D. Jackel and R. A. Buhrmann, J Low Temp Phys, 19, 201 (1975).

JAC75 J. D. Jackson, Classical Electrodynamics, 2nd Ed., New York: John Wiley & Sons, 1975.

JAK83 R. F. Jackiw, Quantum Mechanics, unpublished notes for a graduate course taught at M.I.T.

KUR72 J. Kurkijarvi, Phys Rev B, 6, 832, (1972).

KUR73 J. Kurkijarvi, J Appl Phys, 44, 3729, (1973).

KUW72 J. Kurkijarvi and W. W. Webb, Proc Appl Supercon Conf, Annapolis, IEEE Pub. No. 72CHO 682-5-TABSC (IEEE, New York, 1972), p. 581.

LAL76 Landau and Lifshitz, Mechanics, 3rd Ed., New York: Pergamon Press, 1976.

LAW75 W. N. Lawless, Cryogenics, 15, 273 (1975).

LOU74 O. V. Lounasmaa, Experimental Principles and Methods Below 1 K, New York: Academic Press, 1974.

MCL47 N. W. McLauchlan, Theory and Application of Mathieu Functions, Oxford: Clarendon Press, 1947.

MAR56 J. M. Manley, H. E. Rowe, Proc. IRE, 44, 904 (1956).

MAR70 J. B. Marion, Classical Dynamics of Particles and Systems, 2nd Ed., New York: Academic Press, 1970.

MOO87 F. Moore, Private communication.

NRV87 J. Neukammer, H. Rinneberg, K. Vietzke, *et. al.*, Phys. Rev. Lett., 59, 2947 (1987).

OPS75 A. V. Oppenheim, R. W. Schafer, Digital Signal Processing, Englewood, N.J.: Prentice-Hall, 1975.

OXF85 Oxford Instruments, Superconducting Magnet System manual.

PEN36 F. M. Penning, Physica, 3, 873 (1936). (In German)

PFT86 W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, Numerical Recipes, New York: Cambridge University Press, 1986.

PIE49 T. R. Pierce, Theory and Design of Electron Beams, Princeton: Van Nostrand, (1949).

PRI86 D. E. Pritchard, Atomic and Molecular Physics, unpublished notes for a graduate-level atomic physics course taught at M.I.T.

RAM56 N. Ramsey, Molecular Beams, Oxford: Clarendon Press, 1956.

ROB74 F. N. H. Robinson, Noise and Fluctuations, Oxford: Clarendon Press, 1974.

ROW58 H. E. Rowe, Proc. I.R.E., 46, 850 (1958).

SFG79 W. B. Simmonds, W. A. Fertig, R.P. Giffard, IEEE Trans on Mag, MAG–15, 478 (1979).

SIZ67 A. H. Silver and J. E. Zimmerman, Phys Rev, 157, 317 (1967).

STH51 Sommer, Thomas, and Hipple, Phys. Rev., 82, 697, (1951).

TIN75 M. Tinkham, Introduction to Superconductivity, New York: McGraw-Hill, 1975.

VAS81 R. S. Van Dyck, Jr., P. B. Schwinberg, Phys. Rev. Lett, 47, 395, (1981).

VMF85 R. S. Van Dyck, Jr., F. L. Moore, D. L. Farnham, P. B. Schwinberg, Int. J. Mass. Spec. and Ion Proc., 66, 327, (1985).

VMF86 R. S. Van Dyck, Jr., F. L. Moore, D. L. Farnham, P. B. Schwinberg, Bull APS, 31, 974, (1986).

VOW86 E. Voightman, J. D. Winefordner, Rev. Sci. Instrum., 57, 957 (1986).

VSD87 R. S. Van Dyck, Jr., P. B. Schwinberg, H. G. Dehmelt, Phys. Rev. Lett., 59, 26, (1987).

VWE76 R. S. Van Dyck, Jr., D. J. Wineland, P. B. Schwinberg, H. G. Dehmelt, Appl. Phys. Lett., 28, 446, (1976).

WED73 D. Wineland, P. Ekstrom, and H. Dehmelt, Phys. Rev. Lett., 31, 1279, (1973).

WEL67 P. D. Welch, IEEE Trans on Audio Electro-acoust., AU–15, 70 (1967).

WID75 D. J. Wineland, H. G. Dehmelt, J. Appl. Phys, 46, 919 (1975).

WIL83 M. N. Wilson, Superconducting Magnets, Oxford: Clarendon Press, 1983.

WIN87 D. J. Wineland, Private Communication.

ACKNOWLEDGEMENTS

Everyone who has spent time in graduate school recognizes that a finished, bound dissertation represents one person's reward for the efforts of a much larger group. In my case, this work would never have been completed without the combined efforts of the ICR gang:

—Greg "Kamikaze" Lafyatis, without whose confidence and expertise the experiment never would have gotten off the ground. When he came to MIT four years ago, I couldn't even keep liquid helium in a test dewar.

—Eric "10^{-100} by Tomorrow, or (at worst) the day after" Cornell, whose enthusiasm, quickness, and friendship were particularly vital to our (continuing) success.

—Dave "Big Time, Professional Physicist" Pritchard, whose gentle management of the experiment kept us from turning down too many truly blind alleys, yet still left us feeling that the accomplishments were our own. Obviously, I'm doubly indebted to Professor Pritchard, since he also provided the financial support for my graduate career. (Although neither of us can remember how I ended up here in the first place.)

—Bob "I can't believe I finished!" Flanagan, who demonstrated that it is possible to get through graduate school. Bob performed most of the basic design work for the apparatus, an extremely frustrating and time-consuming task, and despite that work, never had the joy of a successful measurement with the machine.

—Kevin "How can you work without Zappa?" Boyce, whose efforts in the last year relieved me of the onerous task of fixing my earlier electronic mistakes. In addition to his sense of humor, Kevin contributed many excellent pieces of equipment, each of which exceeded both my expectations and my abilities.

By far, the best part of graduate school was my fellow students. Graduate students past—Phil Gould, Peter Martin, Brian Stewart, Pete Magill, Warren Moskowitz—and present—Rick Stoner, Kris Helmerson, Major Dad, David Keith—contributed in many ways to both my education and my well-being. Mike Kash was never to busy to discuss a problem (I can't think of anyone in our lab who doesn't owe Mike some acknowledgement). George Welch and Barbara Hughey, as well as the other Kleppnoids, helped make this place far more bearable.

Finally, I thank Ann Marie for putting up with all of this for so long, even when I wasn't able to explain what the matter was.

If graduate school is like hitting your head against the wall, I can say only that it really does feel much better when you stop!