Improved Single Ion Cyclotron Resonance Mass Spectroscopy

by

Kevin Robert Boyce

A.B. Physics, Princeton 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

October, 1992

Signature of the Author _____
Department of Physics
October 5, 1992

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

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

Improved Single Ion Cyclotron Resonance Mass Spectroscopy

by

Kevin Robert Boyce

Submitted to the Department of Physics
on October 5, 1992, in partial fulfillment of the
requirements for the degree of Doctor of Philosophy

Abstract

We have improved the state of the art for precision mass spectroscopy of a mass doublet to below one part in 10^{10}. By alternately loading single ions into a Penning trap, we have determined the mass ratio $M(CO^+)/M(N_2^+) = 0.999\ 598\ 887\ 74(11)$, an accuracy of 1×10^{-10}. This is a factor of 4 improvement over our previous measurement, and a factor of 10 better than the 1985 atomic mass table adjustment [WAA85a].

We have rebuilt much of our apparatus, increasing the signal-to-noise ratio and improving the reliability of the machine. We have also reduced the typical time needed to make and cool a single ion from about half an hour to under 5 minutes. This was done by a combination of faster ion-making and a much faster procedure for driving out ions of the wrong species.

The improved S/N, in combination with a much better signal processing algorithm to extract the ion phase and frequency from our data, has substantially reduced the time required for the actual measurements. This is important now that the measurement time is a substantial fraction of the cycle time (the time to make a new ion and measure it).

The improvements allow us to make over 30 comparisons in one night, compared to 2 per night previously. This not only improves the statistics, but eliminates the possibility of large non-Gaussian errors due to sudden magnetic field shifts.

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

To Kazuko

CONTENTS

I. INTRODUCTION

The state of the art in precision mass spectroscopy is attained by ion cyclotron resonance (ICR). The MIT ICR experiment in particular has achieved a precision of under one part in 10^{10}, a factor of five more precise than any other measurement [VFS92b]. We expect that simple refinements of our current procedure will yield another factor of two, and with future advances we should be able to get below 10^{11}.

This thesis describes the changes we have made over the past two years, and our new measurement of the N_2^+/CO^+ mass ratio (a factor of four improvement over our previous measurement).

A. Motivation

What can be learned from mass comparisons at parts in 10^{12}? There are several specific physical questions which can be addressed by high-precision mass spectroscopy, such as the neutrino mass and Avogadro's number. There is also the possibility of weighing chemical bond energies. Finally, the accuracy of the atomic mass table can be improved by one to three orders of magnitude even at our current level of precision.

1. The neutrino mass

Several groups are attempting to measure the electron antineutrino mass $m_{\overline{\nu}_e}$ by examining the electrons emitted during the decay of Tritium into ^{3}He, looking at the high-energy end of the energy spectrum [FHK91, KKO91, RBS91]. We can measure the mass (energy) difference between ^{3}He and ^{3}H to better than one eV using one-ion techniques already fully developed. This will allow a significantly improved determination of the upper limit of the neutrino mass.

2. $N_A h$, Avogadro's number times Planck's constant

We can measure the mass difference between mother and daughter species of a gamma-ray decay and combine our result with a measurement of the gamma ray wavelength. The energy of the gamma ray is $hc/\lambda = \Delta mc^2$, whereas the energy of the mass difference we measure is $\Delta Mc^2 = N_A \Delta mc^2$, where h is Planck's constant, c is the speed of light, λ is the wavelength, m is the mass in grams, and M is the mass in amu. Equating the two energies gives $N_A h = \Delta Mc/\lambda$. To improve the limits on $N_A h$ will require precision of better than a part in 10^{11} from our experiment, and an improvement of more than an order of magnitude in the measurement of the γ wavelength. However, even if we cannot improve the limits, this method relies on completely different physics than the previous methods of determining $N_A h$.

3. Weighing chemical bonds

Certain types of molecular ions do not lend themselves to traditional methods of determining binding energy. With precision in the range of parts in 10^{12}, or even 10^{11} in a few cases, we will be able to resolve the existing discrepancies in some of these ions.

4. Atomic mass table

Our measurement of $M(CO)/M(N_2)$ is now nearly a factor of 10 better than the adjusted value from the 1983 atomic mass table [WAH85]. Clearly this kind of spectrometry can make a great contribution toward improving much of the atomic mass table. There are several interesting uncertainties in the mass table at the level of parts in 10^8. Most of these are short-lived nuclei which we can't measure, but some are stable or metastable [AUD91]. Two examples are ^{60}Fe, which has a lifetime of 3×10^5 years, and the doublet ^{76}Ge-^{76}Se, both of which are stable.

B. Review of Penning Traps

In this section, I will briefly review the basics of Penning trap theory, give a short history of our experiment, and list some of the most important measurements that have been made in Penning traps.

1. Some useful theory

A Penning trap is a three-dimensional electromagnetic trap for charged particles. The particles are confined in the radial direction by a constant magnetic field $\mathbf{B} = B_0\hat{\mathbf{z}}$, and in the axial direction by an antisymmetric electrostatic field which provides a restoring force directed toward the center of the trap. For a complete description of the details of Penning traps, see Brown and Gabrielse's review article [BRG86] or Weisskoff's thesis [WEI88]. Both of these references explore the mathematical intricacies of ideal and real Penning traps, so I will provide only an overview here.

The basic idea for making mass comparisons in a Penning trap is that the cyclotron frequency $\omega_c = \frac{eB}{mc}$ depends only on the charge-to-mass ratio e/m, the magnetic field strength B, and the speed of light c. Thus, if we can keep the magnetic field constant (a decidedly nontrivial task, as we shall see), then we can directly compare the masses of different ions using $\frac{m_1}{m_2} = \frac{\omega_{c2}}{\omega_{c1}}$.

As one would expect, the electrostatic field used to confine the ion in the axial direction alters the ion's motion, and we must correct for this before comparing cyclotron frequencies. Fortunately, as we shall see, the free-space cyclotron frequency can be determined from the normal mode frequencies we measure in the trap.

The motion of an ion in the trap can be decomposed into three normal modes: cyclotron, axial, and magnetron. The trap cyclotron frequency ω_c' is due to the usual circular motion in the xy plane around the magnetic field lines, shifted slightly in frequency by the electrostatic field. The axial mode, at the frequency ω_z, is the motion along the z axis due to the restoring force of the electric field. Finally, the magnetron motion, at fre-

quency ω_m, is a slow $\boldsymbol{E} \times \boldsymbol{B}$ drift around the center of the trap in the xy plane. In our trap, the electrode surfaces are carefully machined hyperboloids of rotation (see Figure 1-1), which (ideally) yield a quadrupole electrostatic potential:

$$\Phi(z,\rho) = V_T \frac{z^2 - \frac{1}{2}\rho^2}{2d^2}, \tag{1-1}$$

where V_T is the voltage between ring and endcap, z is the axial position, ρ is the radius, and d^2 is the characteristic size of the trap, defined by

$$d = \sqrt{\frac{\rho_0^2}{4} + \frac{z_0^2}{2}}. \tag{1-2}$$

This results in a harmonic restoring force along the z axis. The frequencies of the three modes of motion in the trap are given by:

$$\omega_z = \sqrt{\frac{qV_T}{md^2}}, \tag{1-3}$$

$$\omega_c' = \frac{1}{2}\left(\omega_c + \sqrt{\omega_c^2 - 2\omega_z^2}\right), \tag{1-4}$$

$$\omega_m = \frac{1}{2}\left(\omega_c - \sqrt{\omega_c^2 - 2\omega_z^2}\right), \tag{1-5}$$

where q is the charge on the ion, m is the mass of the ion, and ω_c is the free-space cyclotron frequency given above as $\omega_c = \frac{eB}{mc}$.

In our trap, the magnetic field is 8.525 Tesla, which gives a cyclotron frequency of around 4.5 MHz for ions of mass 28. The electric field is adjusted to bring the axial motion into resonance with our detector, which is about 160 kHz, and this results in a magnetron frequency of around 2800 Hz for mass 28 ions. For precision mass measurements, we must use the free-space cyclotron frequency ω_c, rather than the modified cyclotron frequency ω_c', since ω_c' depends on the voltage and trap size, neither of which can be measured to the required accuracy. Fortunately, there is an invariance theorem due to

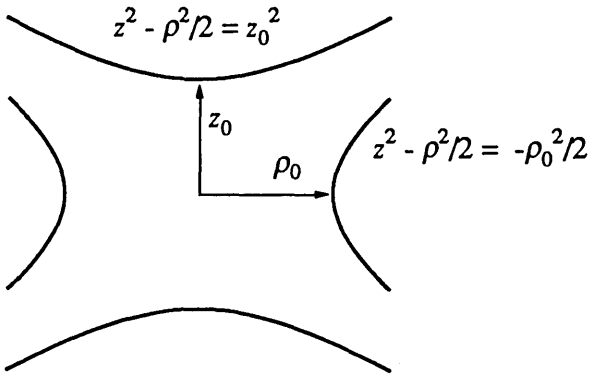

Figure 1-1. The geometry of the Penning trap. The electrodes are hyperboloidal surfaces of rotation. In our traps, $\rho_0 = 0.696$ cm, $z_0 = 0.600$ cm, and thus $d = 0.549$ cm and $d^2 = 0.301$ cm^2.

Brown and Gabrielse [BRG86] which tells us that the free-space cyclotron frequency is the quadrature sum of the three trap frequencies:

$$\omega_c = \sqrt{\omega_c'^2 + \omega_z^2 + \omega_m^2}. \tag{1-6}$$

We always work in a regime where $\omega_c' >> \omega_z >> \omega_m$, so we need the most accuracy for the cyclotron frequency, while the axial frequency requires somewhat less precision, and the magnetron frequency measurement is much less critical.

How do we measure the three normal-mode frequencies? Actually, we directly measure only the axial frequency. We detect the axial motion of the ion by measuring the current it induces in the upper endcap. The endcap is connected to a superconducting tuned circuit ($f_0 \approx 160{,}800$ Hz, $Q \approx 30{,}000$) inductively coupled to a commercial rf SQUID. The theory of our detection scheme was worked out in exquisite detail by Robert Weisskoff in Chapters II and V of his thesis [WEI88]. All his work still applies, although Q and L_1 (the inductance of the tuned circuit) have both been increased (see Section II.F.1).

The cyclotron and magnetron modes are not detected directly. This is to eliminate the perturbations to ω_c' that a radial detector would induce. Instead, we couple the axial and radial modes for short periods of time with an inhomogeneous rf field, allowing us to measure the radial frequencies indirectly. The methods we use for this are reviewed briefly in Section III.B.2, but see also Eric Cornell's thesis [COR90], or [CWB90a], in which these techniques are described in more detail.

2. History of our experiment

The MIT ICR experiment was started in 1983, and first detected ions in the summer of 1986, using our first trap, "Trap 1". We then spent two years trying to see single ions, which we first detected in March 1988. We also installed Trap 2, which had the split guard rings necessary for axial driving and detection, during that time. By early 1989, we had made our first measurement, that of $M(CO^+)/M(N_2^+)$ [CWB89]. Following that we did some systematics checking, measuring $M(N_2^+)/M(N_2^+)$ (which ought to be 1.000...) and $M(N^+)/M(N_2^+)$ (which, after correcting for one electron mass and the binding energy, ought to be 2.000...). In early 1990 we tried loading mass 3 ions, but the vacuum in the trap became too contaminated to use before we could tune the trap properly.

In June 1990 we began the major rebuilding of the experiment described in this thesis. The new apparatus is collectively known as Trap 3. We spent roughly one year building the new apparatus, and another year debugging it and shielding against the increased level of noise in the lab (due to unrelated circumstances). We finally saw a cyclotron resonance in March 1992, and made the new, high-precision measurement of $M(CO^+)/M(N_2^+)$ in May 1992.

3. Some important Penning Trap mass measurements

In recent years, many high-precision mass comparisons have been done in Penning traps, mostly by Van Dyck's group at the University of Washington. Table 1-1, taken

Comparison	Value	Relative accuracy (ppb)	Reference
$m_{\bar{e}}/m_e$	1.000 000 00(13)	130	[SVD81]
$m_{\bar{p}}/m_p$	0.999 999 977(42)	42	[GFO90]
m_p/m_e	1836.152 701(37)	20	[VMF86]
m_p/m_e	1836.152 680(88)	48	[GFO90]
$m_{\bar{p}}/m_e$	1836.152 660(83)	45	[GFO90]
$M(^{12}C^{4+})/4m_p$	2.977 783 713(10)	3.4	[MFS89a]
$M(^{12}C^{4+})/2M(^4He^{2+})$	1.499 161 233(15)	10	[MFS89b]
$M(^{12}C^{4+})/4M(^3He^+)$	0.994 684 341 4(75)	7.5	[MFS89b]
$M(CO^+)/M(N_2^+)$	0.999 598 887 6(4)	0.4	[CWB89]
$M(D_2)$ - $M(^4He)$	0.025 600 331(5) amu	1.25	[GWW90]

Table 1-1 (after [VFS92a]). Summary of some mass comparisons previously performed in Penning traps.

from [VFS92a], lists some recent measurements, including our 1988 N_2^+/CO^+ comparison. Van Dyck has also recently published [VFS92b] [VFS92b]a series of measurements of light ions (H, D, T, ^{3}He, ^{4}He, and ^{16}O), with errors ranging from a few parts per billion (ppb) to 0.5 ppb. These were done by comparing each ion to a "reference" ion, typically multiply-charged carbon.

Penning traps are also being used to determine the masses of short-lived isotopes. Bollen *et al* [BKO92] have measured several francium and radium isotopes, with half-lives as short as 50 seconds, to accuracies of $\sim 1 \times 10^7$.

C. Summary of contents

Here is what you will find in the rest of this thesis.

In Chapter 2, I detail the (substantial) modifications we have made to the apparatus since 1990. We have upgraded (or at least changed) all or part of every section of the experiment, from the trap itself to the data-acquisition system.

Though we aren't quite at the production stage, cranking out measurements weekly, we do have a more-or-less standard set of techniques we use for making mass measure-

ments. I describe the canonical measurement in Chapter 3. This includes the data analysis, with the exception of the first step—estimating the phase and frequency of the ion motion—which is described in Chapter 4.

One of the most important changes we have made is in how we extract the ion's phase, frequency and amplitude from the time-domain signal. Chapter 4 describes what is wrong with our old method and develops our new method.

Chapter 5 presents the most important systematic errors we need to worry about when making a precision measurement, and it gives some estimates of how important each is, both for the measurement we have made and for the upcoming mass 3 measurement. (Statistical errors are dealt with in Chapter 6, along with the measurement data.)

Our best measurement to date is described in Chapter 6. This is, as far as we know, the most precise mass comparison ever done, by a factor of four. The actual data-taking follows the procedures described in Chapter 3, so this chapter is devoted mainly to the analysis of our data.

II. APPARATUS

The experimental setup has been described in detail by Flanagan and Weisskoff [FLA87, WEI88], and additional changes were described in [COR90]. I will give a brief overview of the experimental setup, followed by detailed descriptions of the changes we have made. However, we have changed almost everything, so very little will be missing from the details.

A. Overview of the experiment

The Penning trap is all but buried under a large amount of support equipment, which serves several purposes. First, we need a strong, homogeneous, stable magnetic field. For the ultimate accuracy, the homogeneity must be on the order of a part in 10^8 over a centimeter-sized volume. Temporal stability, especially on time scales of a few minutes, is even more critical; uncorrelated drifts must be below one part in 10^{10}. (This requirement could be relaxed by using two ions simultaneously [CBF92, COR90]).

The DC trap voltage must also be extremely stable, with enough coarse adjustment range for ions from mass 3 to at least 28, and fine adjustment capability below a part in 10^6. The detector must (obviously) be sensitive enough to detect the current of a single ion, and therefore the trap and detector must be exceedingly well-shielded from external fields. This is especially important considering the rather electrically noisy environment in our lab. We need to be able to apply driving and coupling frequencies to the trap electrodes, while keeping out extraneous noise. To minimize Johnson noise, we keep the trap and detector at 4.2 K.

We also need extremely high vacuum. To achieve a precision of parts in 10^{12}, the trap must be below about 10^{-14} Torr. This actually requires little additional effort beyond keeping the trap at 4.2 K. We also need a way to get ions into the trap, without also allowing in noise, causing excessive heat load, or degrading the vacuum.

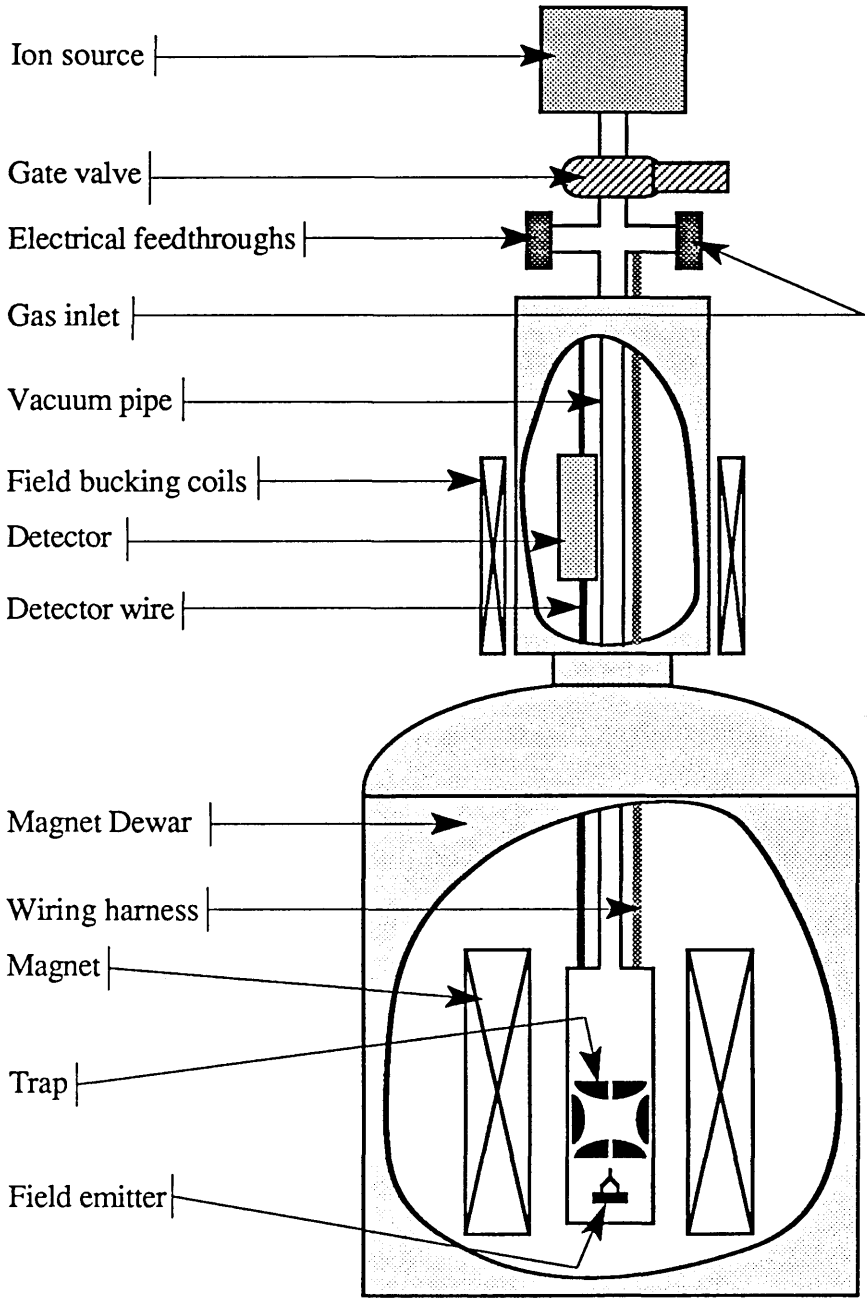

Figure 2-1. The physical setup of the cryogenic apparatus (not to scale). The ion source has not yet been installed.

Figure 2-1 shows the overall setup of the experiment. The trap sits in a copper vacuum can in a liquid-helium filled cryobore at the field center of an 8.5 Tesla superconducting magnet. The copper vacuum can is connected to a copper/stainless-steel tube which connects the trap vacuum to room temperature vacuum, and provides a path for

atoms or ions to reach the trap. A separate copper-and-stainless tube contains all the wires which connect to the trap electrodes.

We make ions in the trap by allowing in a small quantity of neutral gas through a 0.5 mm diameter hole in the center of the top endcap (see Figure 2-2) while also injecting electrons through an identical hole in the bottom endcap. The electrons are produced by a commercial cold-cathode field emission point from FEI, Inc.

We also have provisions for making the ions outside the cryostat and sending them down in a beam. At the top of the room-temperature gas inlet tube is a gate valve, above which we can mount an external ion source. The gas inlet tube contains two Einzel lenses and a deceleration electrode to guide ions into the trap. The ion source is working but has not yet been attached above the gate valve.

The ion signal is detected by a high-Q (~30,000) tuned circuit driving a commercial rf SQUID. The SQUID requires a low magnetic field (< 10 Gauss) to operate, so it is placed in the upper part of the cryostat. (This in fact is the entire point of the upper cryostat extension.) However, the fringing field at the location of the SQUID is about 170 Gauss, so we have wrapped bucking coils around the outside of the cryostat. These allow us to null out the fringing field to a few Gauss, a level which we "freeze" into the superconducting box in which the SQUID is housed.

We have made numerous changes to the entire system during the past year, each of which provides an incremental improvement in the overall performance of the experiment. The insert was rebuilt, using the shell of the insert from the original version of the experiment ("Trap 1"). The trap itself was replaced, as were the cryogenic electronics and all wiring inside the Dewar. We switched to a new SQUID detector and electronics, and we built a new voltage source. We replaced the old LSI-11/23 computer with a Macintosh IIci and added computer control of the ion-making system. We also bought a new signal generator for driving the axial motion and killing bad ions. Finally, we are building an external ion source. Each of these developments will be discussed below.

B. The new trap

Our new trap (see Figure 2-3) is a much more open design than the previous ones. The space between the electrodes is mostly empty, with alumina spacers and alignment posts to support the electrodes. The previous trap was made with solid MACOR rings between the endcap and ring electrodes. (The guard rings were painted onto the MACOR with conductive paint.)

MACOR has two unpleasant properties. First, the small amounts of iron that get incorporated during its manufacture make it slightly ferromagnetic. This results in a distortion of the magnetic field inside the trap, the leading order of which is a B_2 "bottle" term. This term is too large to compensate with the B_2 shim coil built into the magnet, so we added a nickel ring to get within shimming range. This approach worked, but obviously lead to an increase in the higher-order inhomogeneities. We did not calculate the effect of the MACOR and nickel on these higher-order terms.

The other trouble with MACOR was that it was a very lossy dielectric. This is a problem because the endcap-to-ring capacitance appears in parallel with our detector capacitance. Hence some of the dielectric seen by the detector tuned circuit is this lossy MACOR. This (we believe) is what was limiting our Q in our old trap (Trap 2) to 24,000, even though the tuned circuit by itself had a Q of over 50,000. Thus to improve our detection efficiency we had to eliminate the MACOR from the trap design.

The new trap design is shown in Figures 2-2 and 2-3. We had three new traps machined by Ray Harlan, who also made our first two traps. One of the new traps (not shown in the figures) has a split ring electrode. The new traps have guard rings machined from OFHC copper, each supported by 6 alumina step discs. Both guard rings are split along the y axis, allowing us to drive the radial modes of the ion. The lower guard ring is used for driving the magnetron and cyclotron modes, and for mode coupling. The two halves of the upper one are shorted to each other, as in the earlier traps.

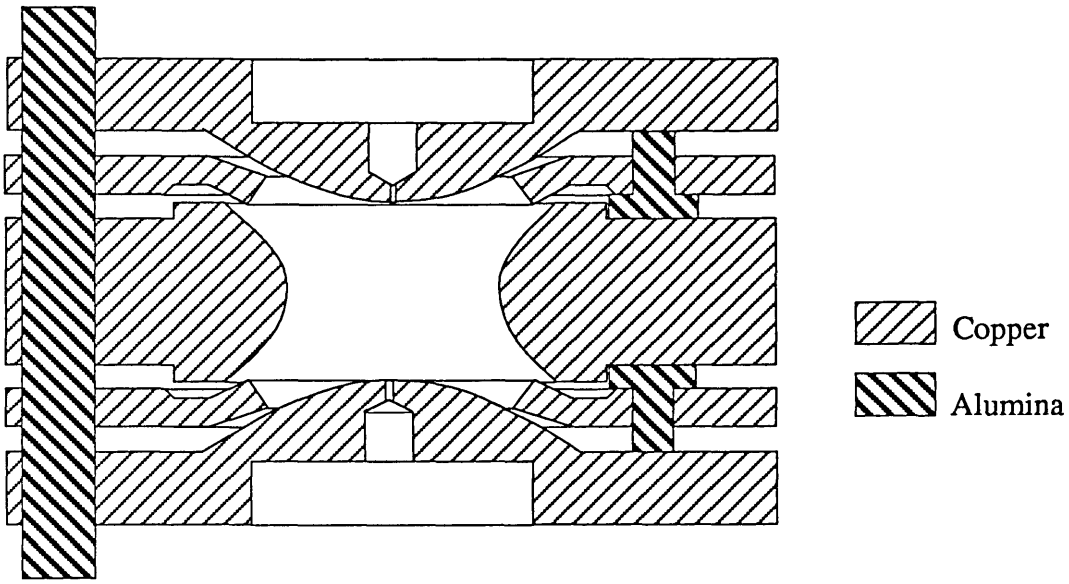

Figure 2-2. Cross-sectional view of the assembled trap.

The large area of the guard rings leads to a substantial amount of capacitive coupling to the ring and endcap electrodes. We measure 7.5 pF between the endcap and the adjacent guard ring. We initially planned to drive the cyclotron motion from the upper guard ring and the magnetron motion from the lower, but it turns out that any circuitry on the upper guard ring limits the Q of the detector (because it is connected to the upper endcap). Both cyclotron and magnetron drives, as well as mode coupling pulses, are applied to the lower guard ring (see "Cryogenic electronics" below).

The trap electrodes are aligned by two alumina rods. The positioning tolerances of the hyperboloidal surfaces were specified to .0003", but we believe that the final product is slightly better than that. Unlike the previous trap, we did not have this one plated, so the machined tolerances have not been compromised. In fact, the effect of geometric effects on electric field inhomogeneity is negligible compared to the effect of surface charge patches.

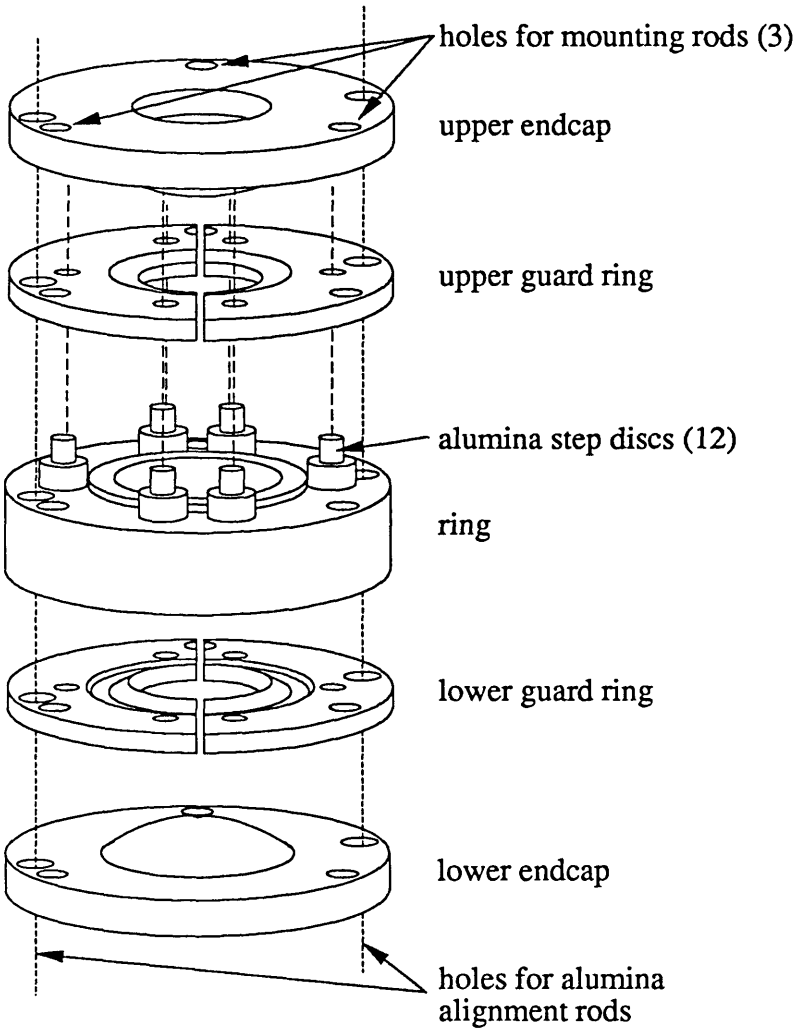

Figure 2-3. Exploded view of the new trap. Each step disc has a copper-beryllium spring around it (not shown), between the endcap and the guard ring, to hold the guard rings in place against the step. There is a small Teflon washer under each spring, for electrical insulation. There are six more alumina step discs for the lower endcap, which are not shown.

We used the following procedure to minimize charge patches. All the copper surfaces were cleaned with a dilute acid bath followed by deionized water, methanol acetone and methanol again, and finally sprayed with Aerodag G. Aerodag is a coating of carbon particles 10 microns and smaller in a carrier of isopropyl alcohol, in a convenient aerosol can. According to Camp et al at Los Alamos, [CDB91], Aerodag on clean copper surfaces is the best way to minimize surface potentials. One problem with Aerodag is

that it has a rather high resistivity at 4.2 K, which caused a dramatic drop in the Q (to about 14,000—more than a factor of 2). However, this problem was easily solved by removing all the carbon from the facing surfaces of the upper endcap and guard ring (Aerodag is easily removed with alcohol).

The Aerodag will not substantially reduce the patch effects if it has lots of contaminants frozen onto the surface, so we tried to be more careful than in the past to keep the vacuum can clean. To this end, we removed all the rosin flux from the cryogenic electronics with acetone before initially installing the trap. (After the *N*th time of rewiring the cryoelectronics, we were somewhat less scrupulous, however. A complete cleaning would probably be helpful the next time the system is cycled to room temperature.) We were also much more careful not to get skin oil on the electronics or trap.

The coiled copper cryoadsorber [COR90] was replaced, at the suggestion of Steve Jefferts (of JILA), with an activated carbon adsorber. This consists of activated charcoal pieces, about 0.5 cm diameter, glued to copper screening with Torr-Seal. The copper is heat-sunk to the copper vacuum can with Cu-Be fingerstock. This, combined with the much more open design of the new trap, should improve the vacuum when using light ions (H and He), which have high vapor pressures at cryogenic temperatures.

C. Voltage box

The voltage source for the trap must be extremely stable, since $\Delta\omega_z/\omega_z$ is just $\frac{1}{2}(\Delta V_T/V_T)$. Recalling Equation 1-5, we see that the fractional accuracy required for the axial frequency is less than that required for the cyclotron frequency by the factor $(\omega_z/\omega_c')^2$. For example, if we want a final accuracy of a part in 10^{11} using mass 18 (where ω_z/ω_c' is 1/50), we require $\Delta V_T/V_T < 5 \times 10^{-8}$. This is beyond the limit of the previous voltage source. Also, the higher Q of our new detector necessitates improved stability of ω_z, because the pulling of ω_z by the detector resonance is dependent on the

detuning. Thus variations in the axial frequency result in variations in the pulling of the axial frequency. For these reasons, we built a new, more stable, voltage source.

A simplified schematic for the whole box is shown in Figure 2-4. The major source of drift in both the old and new voltage sources is thermal coefficients of various components. Several parts of the old supply contributed at the 5 to 10 ppm/°C level: the voltage reference was an LT1021 (specified at <5 ppm/°C); the output offset of the AMP-01E instrumentation amp (which was used to add the computer-controlled offset to the reference) is about 50 ppm/°C, (this was reduced by an 11:1 resistive divider); and the voltage adjustment potentiometers (Clarostat 62JA) have division ratios which vary between 1 and 10 ppm/°C, depending on the temperature and the exact position of the wiper.

The new version of the voltage box uses the Linear Technology LTZ1000 voltage reference, which claims a stability of 0.05 ppm/°C, two orders of magnitude better than our old reference! Also, the divide-by-10, 100, or 1000 networks which were between the computer and the voltage box have been placed after the instrumentation amp, to reduce its effect on V_t to around 0.05 ppm/°C. Finally, the potentiometers are now used as a fine adjustment to a 4-bit hand-switched D/A which multiplies the output of the reference chip. All the voltage-setting resistors, including those in the D/A network, are Vishay™ metal foil types, with temperature coefficients between 0.2 and 1.5 ppm/°C. All in all, the new voltage box shows a temperature coefficient of about 0.5 ppm/°C, and it can easily be controlled to within about 0.1°C, so we can stabilize V_t to better than a part in 10^7 over times of an hour or so.

We also added several convenience features, such as the abilities to select between two different voltages, apply high voltage for a high-frequency cooling scheme, and dip the ions toward one endcap.

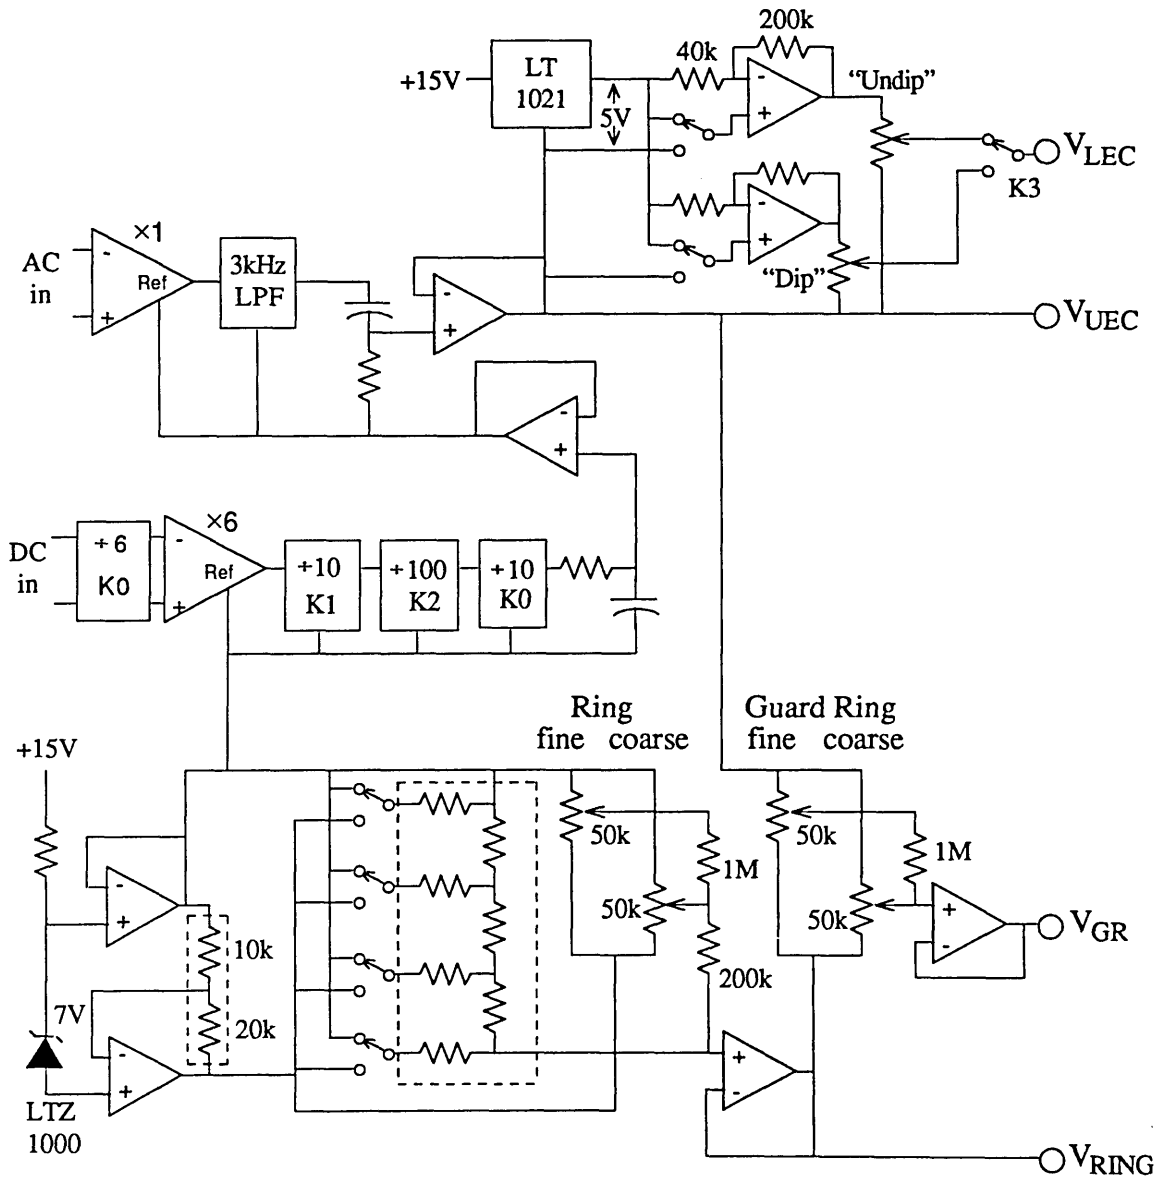

Figure 2-4. Simplified schematic of one channel of the new voltage box. There is duplicate circuitry (which is selected by a set of relays) for the ring, guard ring, and dip voltages.

D. Cryogenic electronics

All the filters and driving electronics which live in the copper can right above the trap were rebuilt in a way to make them less likely to break (due to thermal cycling or physical stress when servicing them), but are electrically very similar to what they were previously. The circuitry is shown in Figure 2-5.

The DC trap voltages are filtered by 4-pole LC filters with a cutoff frequency of about 300 Hz, except for the ring, which has a cutoff of about 2 kHz. All excitation is applied through transformers. The transformers and inductors are wound of 34 AWG copper wire with heavy-duty insulation (Belden "Heavy Poly-Thermaleze™")on Teflon forms. Each coil has a small piece of G10 circuit board on the end, with terminals for connecting to other components. This is much more reliable than the old method of simply soldering the coil wire to another wire and wrapping it with Teflon insulation.

The axial excitation transfer function has a gentle ($Q \approx 1$) peak at about 160 kHz, and falls off above and below that. The cyclotron transfer function is less than -80 dB below 500 kHz, and has a peak around 4.5MHz with a Q of about 5. This is due to the self-resonance of the transformer, and is undesirable but not impossible to live with. The magnetron transfer function, oddly, has its peak at about 20 kHz, while the frequencies of interest are near the axial (160 kHz) and magnetron (200-3000 Hz) resonances. This is again not ideal, but a compromise among passing all the frequencies desired, circuit complexity, and design time.

One serious problem with this design is the 1 kΩ input impedance of the cyclotron drive transformer. This of course causes a very large standing wave ratio in the driving cable, leading to rather sharp resonances in the cyclotron drive transfer function. In fact, the cable length from the drive electronics to the experiment puts the 4.5MHz frequency of mass 28 ions on the steepest part of the first resonance. This is another problem which

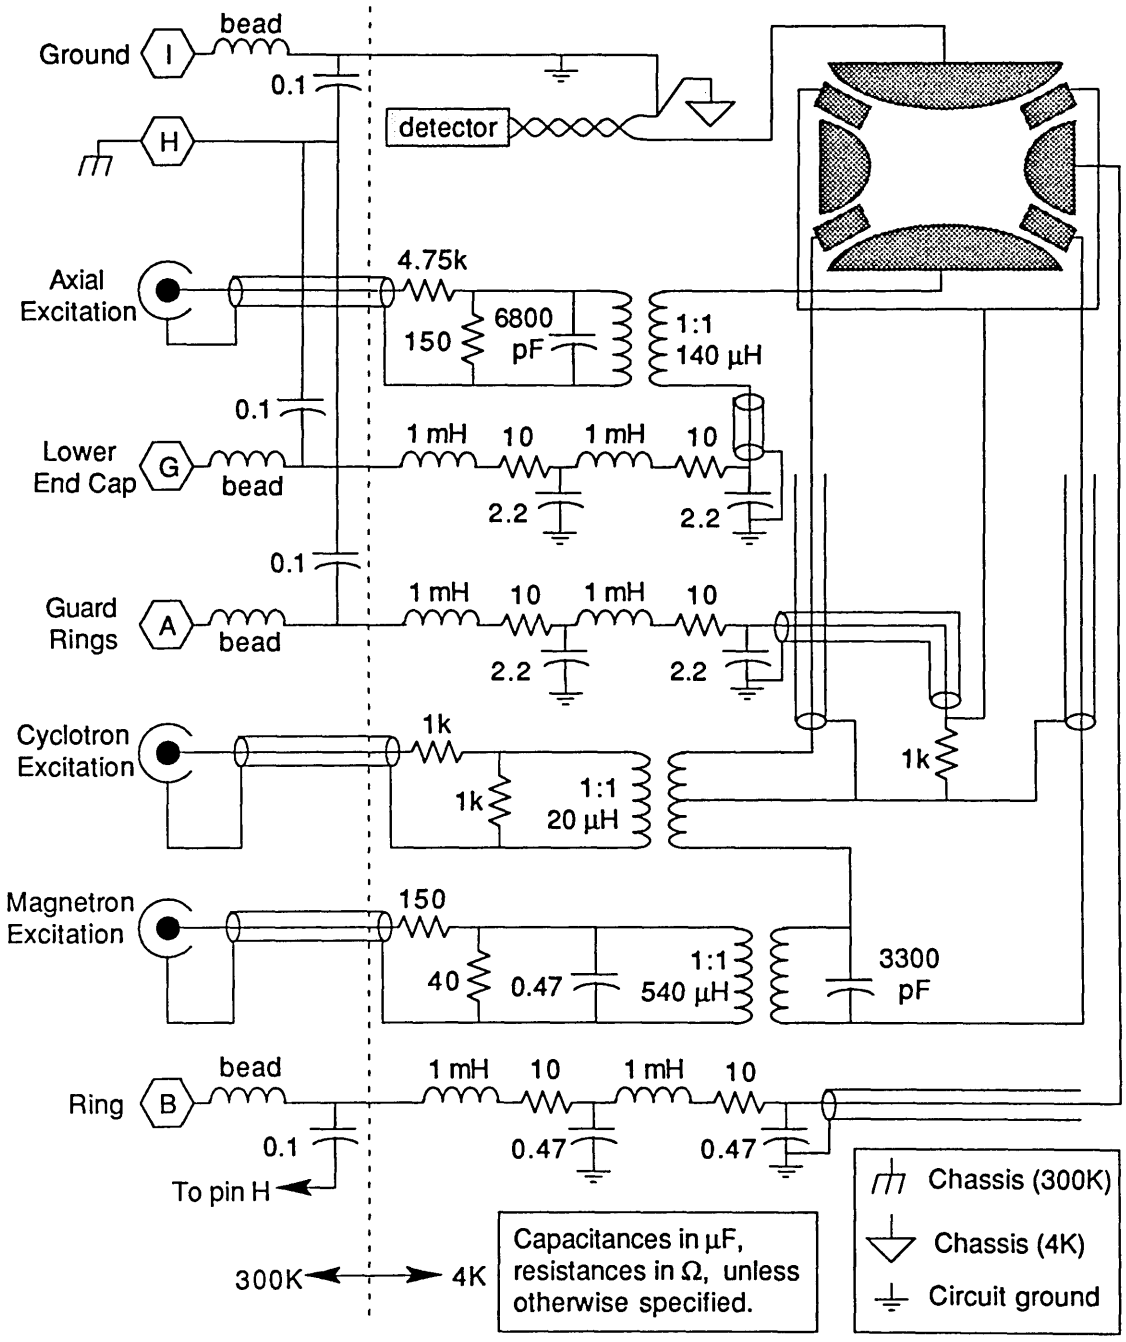

Figure 2-5. The cryogenic filtering and coupling circuits. The "bead" inductors are ferrite beads (Fair-Rite material 73) on the feedthrough leads. There is approximately 2 meters of wire (subminiature 50Ω coax[1] or 5 mil copper wire coated with Teflon) between room temperature and the cryofilters.

[1]Type C1 coax, from RMC Cryosystems, 1802 W. Grant Rd., Suite 122, Tucson, AZ 85745

needs to be fixed the next time the insert is pulled out of the magnet. In the meantime, we have added enough cable to put the first resonance peak at about 4.5MHz, so that the response is approximately flat for various mass 28 ions. We also plan on adding a 1:6 step-up transformer (impedance ratio 1:36 = 50:1800) at the room-temperature end. This will provide slightly better impedance matching, although there is 2 meters of 50Ω coax between there and the cryofilters.

When the new trap was first put together, the detector was connected to the lower endcap, so that the variable voltage for shifting the ions vertically in the trap was applied to the upper end cap. The reason for this was so that the upper endcap could now easily be used as a gate to trap incoming ions. Unfortunately, the lower endcap is also very close to the Field Emission Point (FEP), which apparently limited the Q of the detector to about 25,000 (probably due to the Aerodag on the endcap). Also, the resonant frequency of the detector shifted by 50 Hz when we plugged in the room-temperature high-voltage supply to the FEP. Thus we returned the endcaps to their original roles, with the detector on the upper endcap, and the "dip" voltage on the lower.

The axial drive is applied to the lower endcap. The magnetron and cyclotron drives are combined and applied to the lower guard ring. We had to modify the guard ring circuitry slightly from what it was before. The large area of copper guard ring in the new trap has a large capacitance to the endcap (7.5 pF from one guard ring to one endcap). This is a substantial fraction of the total capacitance of our detector. Thus it is critical, to maintain a high Q, that this is "good" capacitance (i.e. not lossy). Thus the upper guard ring, which is the one with a strong coupling to the detector, is connected directly to ground through a large (2.2 μF) capacitor.

E. Driving the ions

As will become clear in Chapter 3, we need to be able to apply coherent pulses and continuous drives of various frequencies to the trap. For the axial motion, we drive the

ion in both pulsed and CW modes, and we apply white noise to drive out ("kill") ions of other species. For the radial modes, we apply drives to excite the cyclotron motion, and to couple magnetron or cyclotron modes to the axial mode. When actually measuring the cyclotron frequency, we apply a pulse of cyclotron driving frequency followed by a pulse for cyclotron-axial coupling. It is important to maintain a fixed phase relationship between the two frequencies over the course of a measurement.

We have revamped the system for driving the various modes of the ions, with a new axial drive signal generator, notch filter to shape the white noise for killing, and a real coax relay for multiplexing the cyclotron coupling and driving frequencies. Figure 2-6. shows a block diagram of the axial and radial driving electronics.

1. Axial drive

We now drive the axial mode with a Stanford Research Systems (SRS) DS345 signal generator, which has several very nice features for driving and killing ions. It has a burst mode, in which a given number of cycles is generated, starting at a trigger. This allows us to eliminate the Millisecond Pulser which we used to use for pulsing the axial motion. Generating the pulses directly with the SRS burst mode is somewhat more flexible than using the pulser. In addition, we eliminate the extra noise from the pulser, and in the "off" mode there is zero feedthrough, since the oscillator is not even running (the SRS is a direct-digital synthesis unit).

2. Killing

After making an ion, we generally also have a few ions of other species, which we call bad ions, in the trap. They cause uncalibrated and time-varying shifts in the axial frequency and anharmonicity of the "good" ion, and thus must be eliminated. We do this by driving with shaped white noise and dipping the resulting excited cloud of ions close to the lower endcap (see Chapter 3 for more explanation).

The noise is made by the SRS, which can generate pseudorandom white noise which is flat to 10 MHz. We use this noise, after passing it through a passive LC notch filter centered at the detector frequency, to excite the bad ions. This has reduced our ion-killing time from about 12 minutes to under one minute. In addition, since it comes from the same generator that makes our axial drive signal, we do not need to change cables (or add a multiplexer) to switch between driving and killing.

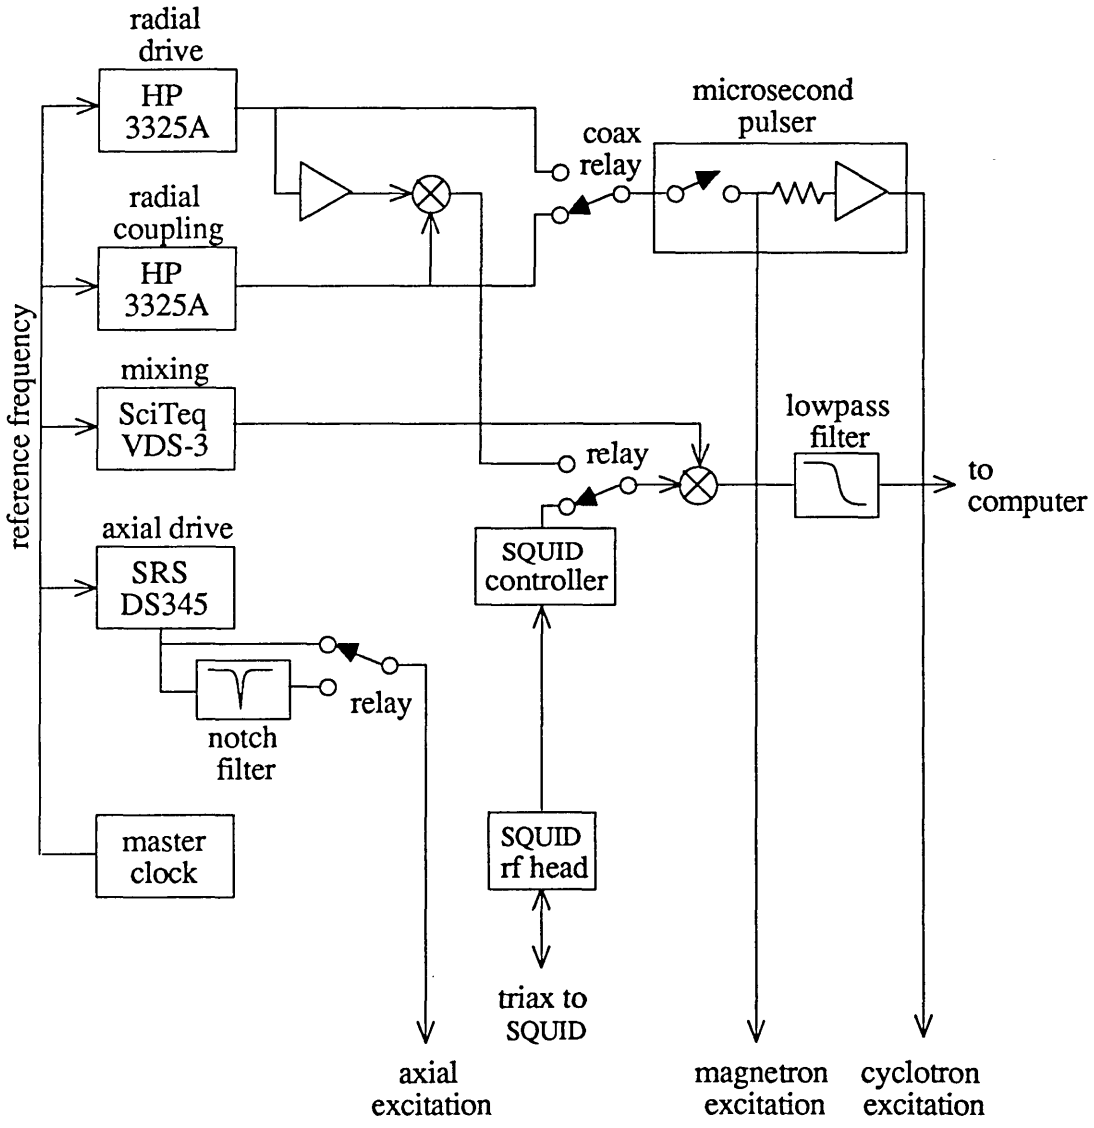

Figure 2-6. The driving and detecting electronics. The relays are all controlled by the computer, as are all the signal generators except the SciTeq. Also, the SQUID can be put into reset mode by the computer.

The notch filter has a -3 dB width of about 40 kHz and a depth of about -30 dB. This is about a factor of two wider and a factor of 10 (20 dB) shallower than we would like, but we probably need an active filter to get that kind of response. Since we need to pass signals up to about 1 MHz, we can't use ordinary op-amps, so we have not yet taken the time to build an active version.

The SRS also has a provision for downloading an arbitrary waveform of up to 16,000 points. This opens up the possibility of generating digital noise with a very sharp notch from a high-order digital filter. This may be useful for light (mass 3) ions, since when they are tuned to the detector frequency the heavy ions (gold, tungsten, etc.) are down around 20 kHz, where the transfer function of the cryofilters is falling off pretty rapidly. Thus it is hard to get enough drive to excite the heavy (bad) ions sufficiently without also driving the good ion out of the trap.

3. Radial drive

Cyclotron and magnetron signals (driving and coupling) are still provided by our two HP 3325A synthesizers and our "microsecond pulser", but we have improved the arrangement for switching between driving and coupling frequencies. This was previously done with a general-purpose relay, which provided reasonable isolation for mass 28 ions (~4.5MHz), but was questionable for mass 14, and unacceptable for mass 3 (~43 MHz). We now have a coax relay rated at 60 dB of isolation up to several hundred MHz.

In addition, we have added an amplifier and mixer to determine the relative phase of the driving and coupling pulses. The reasons for this will be described in Section III.B.2. The input signal to the computer can be switched between the SQUID output and the output of this mixer with a relay.

F. Detector

The detector is fundamentally unchanged, a tuned circuit feeding an rf SQUID, but the SQUID has been replaced, and the tuned circuit replaced with one which has different parameters. The guts of the detector are shown in Figure 2-7.

1. Tuned Circuit

We have increased the coil inductance of the tuned circuit from about 5mH to 9mH. This reduced the uninstalled Q to about 50,000 (we got up to almost 80,000 with the 5mH coil), but that is still twice the installed Q of the old detector. To keep the frequency up near 160 kHz, we reduced the capacitance accordingly. In fact, there is now no superconducting capacitor at all; the capacitance (about 105 pF) is entirely parasitic. There is about 5 pF of self-capacitance in the coil, 40 pF between the wires leading to the trap, and 8 pF from the upper end cap to the guard ring. The remaining 50 pF or so is between the "hot" wire of the twisted pair and the (grounded) copper shield around it .

2. SQUID

The SQUID sensor, rf head, and control electronics have been replaced with units from Quantum Design, Inc[2]. This system is substantially quieter than the old one, along with being cheaper and more rugged. In addition, the previous manufacturer, BTI Inc., no longer makes or supports SQUID current probes.

3. Shielding and Q

The old detector was contained in a niobium box, to provide shielding from both electromagnetic interference and magnetic noise. Combined with a watertight layer of lead foil (COR90), that scheme was quite effective. However, the mechanical arrange-

---

[2]Model 2000 rf head, model 2100 control unit, thin-film rf sensor, from Quantum Design, Inc., 11578 Sorrento Valley Rd., San Diego, CA 92121

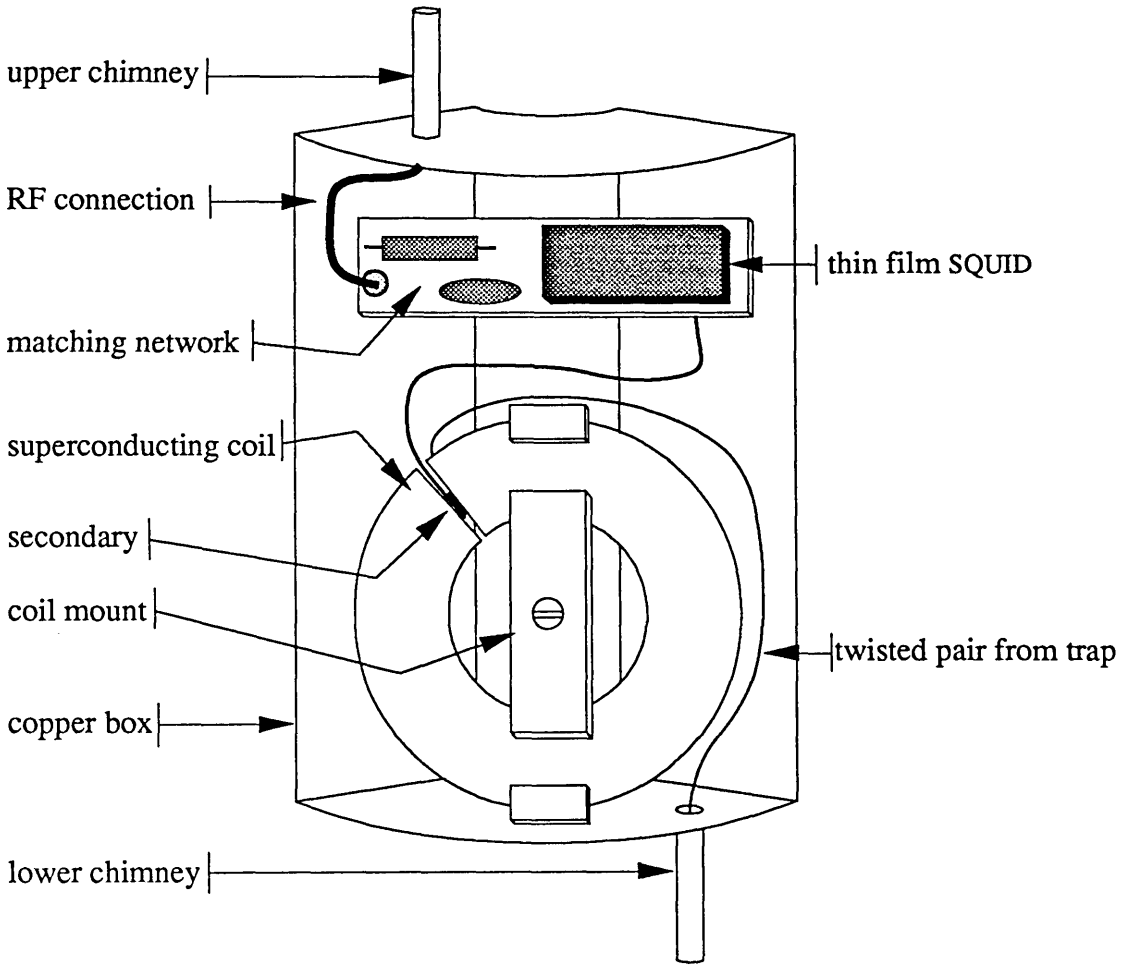

Figure 2-7. The detector box, with its cover removed. The inside surfaces of the copper box are coated with lead.

ment required very careful assembly to avoid pinching the SQUID input wires when closing up the niobium box.

Our new detector box is front-loading, allowing the SQUID, coil, primary and secondary wires to be attached easily with no risk of pinching or stressing anything. The whole assembly, as shown in Figure 2-7, is then covered with a loose-fitting copper lid. The magnetic and electromagnetic shielding is provided by lead foil, as before, but we found that one layer of 8 mil foil did not provide sufficient magnetic shielding.

Our first attempt to increase the shielding was to wrap the SQUID sensor in its own lead shield within the outer lead bag. For reasons we still don't understand, this reduced the Q drastically. We measured the Q and magnetic shielding in our test Dewar, using a 4" permanent magnet from an ion pump to check field penetration. When the lead bags were both tightly sealed with solder, placing the magnet against the outside of the Dewar (about 15 cm from the detector) caused a change of less than 10 flux quanta through the SQUID loop. The measured Q, however, was always around 200. With the outer bag tightly squeezed shut but not soldered, we got a Q of 3-4000 and moderate shielding of the DC magnetic field. Finally, with both bags sealed, we could pull the detector out of the Dewar, quickly slit the outer bag vertically and drop it back into the liquid helium. The Q would invariably go back up to around 50,000, while the magnetic shielding factor dropped by 3 or 4 orders of magnitude.

We do not understand this behavior, but it is clearly related to the bag-within-a-bag topology. Perhaps the flux transformer (the secondary circuit) coupling flux from between the bags to inside the inner one causes enough physical force to flex the lead. Lead is very soft and very mechanically lossy, so this seems plausible. In any case, the solution was to put both lead bags around the outside of the detector. The inner bag is soldered to the shields of both the wires to the trap and the rf triax to the room-temperature electronics. For the Q to be high, the outer bag must have no superconducting (lead to lead) connection with the inner one, although normal conductivity (through the stainless steel support tube) is not a problem.

The points where the rf triax and input wires penetrate the lead bag are potential places for magnetic field to penetrate. To minimize this penetration, they enter through long (~5 cm), thin (~2 mm dia.) "chimney stacks" of lead, which are soldered to the inner lead bag. Magnetic field penetration into a hollow tube goes as $e^{-1.8l/d}$, where l is length and d is width [RMC90], so these chimneys presumably provide a field attenuation of e^{-45}.

In actual use, the new Q varies from 25,000 to 35,000 from day to day. Letting the liquid helium level drop to near the bottom of the lead bag and refilling it usually changes the Q by 10-20%, so we assume the problem is in the detector rather than the trap. Also, because of the unusual behavior described above, we suspect the lead bag is the cause of the variation of Q.

We are currently designing a new detector can, using all-niobium technology, but designed so that the circuitry can be assembled separately and then slid into the can. In order to minimize magnetic field penetration, we need to have complete (watertight) coverage laterally. That is, we need a superconducting path completely surrounding the detector in the x-y plane. Thus the geometry must be that of a can, open only on the top, with a tight-fitting lid that extends far down the sides of the can. If we keep the gap between cover and can to less than 20 mils, a 2" overlap will allow only e^{-100} penetration of the field. Similarly, we will continue to use the present chimney stack scheme where the input wires and rf coax enter the box, although the chimneys will be of niobium, rather than lead. We will also attach copper bushings to the ends of the chimneys, probably by shrink-fitting, to allow easy soldering to the outer shields of the triax and twisted pair.

4. Bucking coils

The Quantum Design SQUID sensor is apparently more sensitive to magnetic fields than the BTI sensor was; it will not operate in the ~170 Gauss field which exists at its location. Therefore we have to null out this field somehow. In theory, lead or niobium, being type I superconductors, should expel all the field lines in them when they make the superconducting transition (Meissner effect). However, in practice there are defects which allow flux to penetrate, and the effect of the transition is to freeze in whatever field exists.

Thus we need to buck out the fringing field while we fill the cryobore with liquid helium. We do this with a set of coils wound around the outside of the Dewar (see Figure

2-1). There is room for one layer of 14 gauge square magnet wire between the insert Dewar and the "towers" for filling the magnet Dewar [FLA87]. However, that is not enough to generate the necessary field without exceeding the power rating of the wire, so we added a layer of 12 gauge round wire outside the towers. This creates the necessary field at the position of the SQUID, though it takes 31 Amps and dissipates over 500 Watts when warm. It's a good thing the lead bag holds the field, so we can turn the bucking coils off once it becomes superconducting.

G. The new computer

We have replaced our old LSI-11/23 computer with a Macintosh IIci. This gives us much more processing power, which we need for our new data analysis scheme (see chapter 4). Each ring-down now takes about 15 seconds to process; it would take about 2 minutes on our old machine. The additional power also eliminates many of the compromises we had to make to see the data in real time.

For example, we used to take almost all our data using a digital filter and 4x downsampling [WEI88], primarily to keep the data array from exceeding 1024 points. This filter introduced phase and amplitude errors which, though small, were still significant. We can now take all our data directly and "filter" in the frequency domain by simply ignoring the higher frequency components.

Another improvement is that the DACs in the new system (National Instruments NB-MIO-16) are far more accurate than the old ones. We now get agreement between the DAC and ADC of better than 1/4 of a least significant bit (LSB) at any voltage. This is important for slow sweeps in the lockin mode, where we might be changing the voltage by one LSB per second. In particular, there was a 10 mV (2 LSB) discontinuity at 0V with the old system; that is no longer the case.

There are also many convenience features that we didn't have before, such as more data storage, automatic data backup, and an excellent analysis and graphing package[3].

The front end of the data-taking software is LabView™[4], a graphical programming language designed for acquisition and control. It provides a reasonable user interface and is convenient for building a new "Virtual Instrument" out of existing modules, but is exceedingly slow for anything other than simple data-taking. Fortunately, LabView provides a reasonably convenient mechanism for including C code in a program, so we have written all the time-sensitive routines (pulsing, lockin amplifier, *etc.*) in C.

H. The magnet

The magnetic field of our trap is provided by an Oxford Instruments 360/89 superconducting NMR magnet. When initially shimmed by the manufacturer, it had a homogeneity of 1.4×10^{-8} over a volume of 1 cm^3[FLA87]. The permeability of the trap reduces that value somewhat. In addition, the shim coils are not optimally set right now, for the following reason.

In December of 1990 we had to discharge the magnet, and we recharged[5] it ourselves. We had brought a large power supply too close to the experiment, causing the magnet to shift in the Dewar, crushing the superinsulation and dramatically increasing the boiloff rate. (The official term for this is a "light touch".) We discharged the magnet (quenched it while attempting to discharge it, actually), and found that once the field was gone the magnet apparently relaxed back toward the center, reducing the boiloff to the normal rate. The magnet may still be off-center, leading to a large radial asymmetry in the magnetic field (see Section V.C.1).

---

[3]Igor, by WaveMetrics, 10200 SW Nimbus Ave., #67, Portland, Oregon 97223
[4]National Instruments, XXX, Houston, Texas XXXXX
[5]We always say "discharge" and "recharge," although "recurrent" would be more accurate. What we do is ramp up the current to the appropriate value and switch the magnet back to persistent mode.

Since we have no way of measuring the inhomogeneity ourselves, we simply reset the shims to their previous values, within about 0.2%. However, we could only set the main field to about 3% accuracy. When we measured the cyclotron frequency, we found that the field is about 1% higher than before the quench, so the shims are off by 1% of their value. Most of the shims have a range of about a part in 10^6 over a cm^3 volume, and most are set to one quarter to one half of their range, so the maximum resulting inhomogeneity is on the order of a few parts in 10^8, which is close to the original spec. However, the fact that the magnet physically shifted in the Dewar means that the field center may not be at the center of the cryobore.

If the magnet is no longer centered, the field homogeneity may be substantially degraded, particularly the asymmetric radial components (X, X^2, XY, etc.). In addition, the axial linear (Z_1) and bottle (Z_2) shims will have a substantial asymmetric component. Thus, even if the homogeneity is initially acceptable, when we adjust the axial shims to compensate for the presence of the trap material we will introduce an asymmetric radial term. Though such a term will be averaged out by the cyclotron motion, any voltage-dependent offset in the radial position (due to a charge patch, for example) will result in a systematic shift of cyclotron frequency with trap voltage. We attribute the error in our atomic-vs-molecular N_2 measurement to this effect (see Section V.D).

Also, the field now points down, whereas before it pointed up (the current in all the shim coils has also been reversed). Assuming that the magnetic materials near the experiment are reasonably linear, this should not have any effect on the experiment (until we can see the effect of the earth's rotation—about 5×10^{-12} for mass 28—and even then, since we can't compare it to anything, it will still be undetectable).

I. The insert

The trap, SQUID, and cryogenic electronics are all mounted on a 2 meter long insert which can be removed from the magnet bore for modification. (See Figure 2-8.) The in-

sert consists of the copper vacuum can, a main tube for loading the trap, and a smaller tube (also connected to the trap vacuum) containing wires to the trap. There is also a box containing our detector, and wires from it down to the trap and up to the room-temperature SQUID electronics. The detector and its wires are immersed in liquid helium.

Near the top of the insert is a radiation shield which is thermally connected to the 77K walls of the cryostat by Cu-Be fingerstock. Just below that is a set of spiral baffles, designed to maximize the cooling effect of the escaping helium gas. Three centimeters below the baffles, the cryobore widens to provide a 4.4 liter reservoir for helium.

Near the bottom of the helium reservoir is the SQUID and detector. The connection to room-temperature SQUID electronics is via miniature coax inside a cupronickel tube. A twisted-pair wire inside a copper tube runs from the detector down to a homemade feedthrough into the copper vacuum can and thence to the trap.

When we decided to add the capability to make ions externally, we wanted to maximize the clear path down the central tube, in order to minimize our ion-steering requirements. We calculate that, at an energy of around 100 eV, the ions will be caught by the magnetic field about 50 cm above the trap. Then, since the field lines are converging toward the trap, any ions that we can get that far down will be "funneled" into the hole in the upper endcap by the field. Thus the condition for getting ions in is just to avoid hitting the beam tube before getting to the region where the magnetic field begins to guide them.

To maximize the size of the beam line, we added a second, smaller, tube off to the side. It has a diameter of 3/8", and carries all the DC potentials and AC excitation to the trap. It is part of the main vacuum system, so no additional feedthroughs are created by this scheme. Both the wiring harness and main tube are copper near the trap (to minimize magnetic field gradients) and stainless steel in the neck of the cryostat (to minimize heat transfer).

To reduce heat flow (and the liquid Helium boiloff it causes) we used much thinner (.005" diameter) copper wire for the DC voltages, and replaced the coax wires for excitation with stainless steel coax. These two changes have reduced the helium boiloff rate of the insert by a factor of two, to about two liters per day. This not only cuts our cost for cryogens nearly in half, but also reduces any time-dependent effects caused by the change in liquid helium level over the course of a measurement.

Figure 2-8 shows the relative locations of the important parts of the insert, their depths below the top plate, and how they relate to the bore of the cryostat. Since the thermal contraction is substantial, it is important to note that these measurements are made while the insert is warm, . When we first measured B_1 and B_2 (see Section V.C.1), we found that the B_2 shim coil had a large B_1 component, consistent with the trap being about 8 mm above the field center. We then remembered that shrinkage was a problem, and we estimated that the change in length from the top plate to the trap center when we cooled the insert to 4.2K would be about 6 mm. Fortunately, all the tubes which pierce the top plate are anchored by Cajun fittings and fixed vertically by aluminum spacers above the plate. Therefore we were able to lower the trap by 6-8 mm by simply machining a shorter spacer.

The number shown in Figure 2-8 for the trap location, -165.0 cm, is correct. This number comes from the magnet manual, and is independent of the temperature of the bore, since the top plate is supported by room-temperature material. The magnet itself rests on the bottom plate of its Dewar, and the top plate is supported by the outside of the stainless steel Dewar. The cryobore hangs down from the top plate, and does not touch the bottom of the magnet Dewar. This leads to undesirable motion of the trap in the magnetic field with temperature variations in the lab, but at least there is no ambiguity in the location of the field center with respect to our reference surface, the top plate.

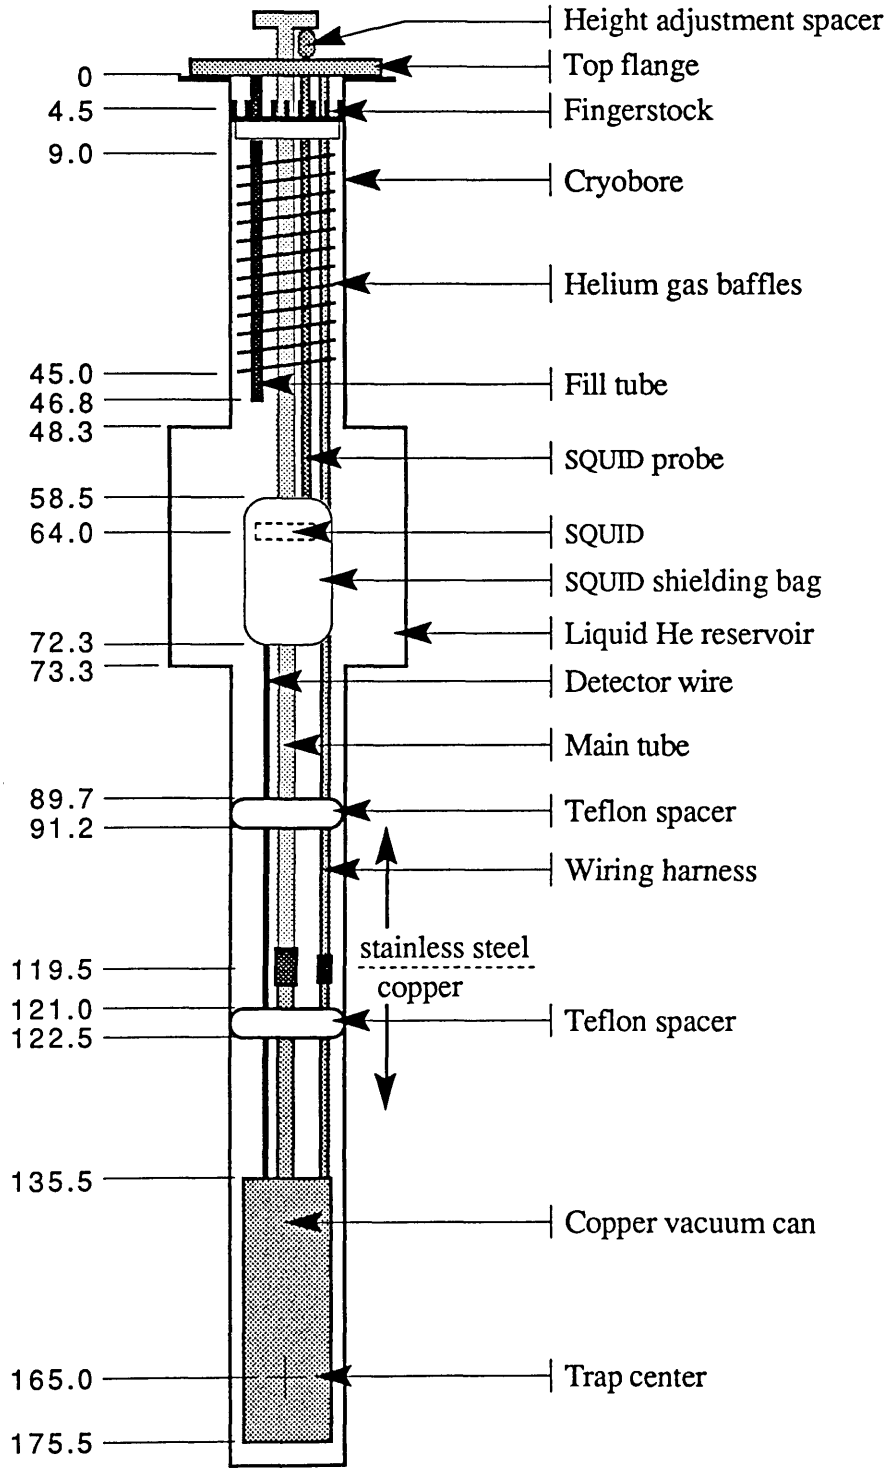

Figure 2-8. The cryogenic insert in the cryobore (not to scale). Numbers on the left are depth below the bottom surface of the top flange, in centimeters, measured at room temperature.

J. The new gas handler

The gas handling system has been completely rebuilt, with all welded stainless-steel connections in the critical locations, a turbomolecular pump, and computer-controlled valves. A diagram of the system is shown in Figure 2-9.

The main section of the gas handler (two storage bottles, expansion chamber, and

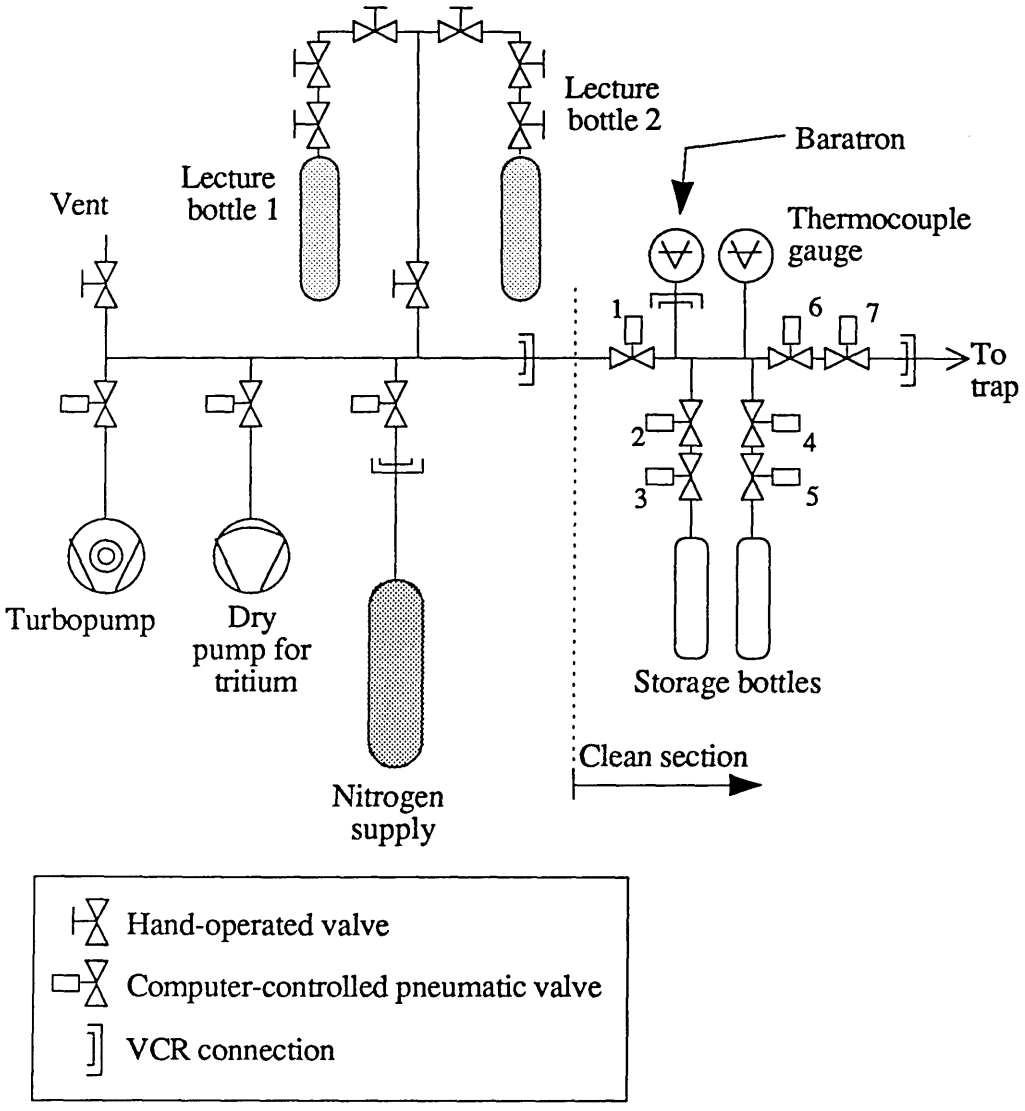

Figure 2-9. The major parts of the gas-handling system. The clean section is stainless steel, all welded except for the VCR connection to the Baratron and the Swagelok and pipe-thread connection to the TC gauge. Swagelok fittings are not shown.

metering chamber) is all welded except for a Swagelok and 1/4" pipe thread joint to the thermocouple gauge and VCR fitting to the Baratron. The valves are still bellows-sealed high-vacuum types with KEL-F stem tips (Nupro SS-4BK), but we have replaced the manual stems with pneumatic actuators, to allow automated ion-making. The valves are all controllable from our computer, and the most commonly used ones are also controllable with a small switch panel near the gas handler.

Our new ion-making scheme (see Section III.B.1) is efficient enough that we typically run with 5-15 mT of N_2 or slightly more of HD in the main expansion chamber. The mechanical pump we used previously had a base pressure of maybe 10 mT if the gas handler was especially clean, so obviously that pump is no longer acceptable. We have built a pumping station based on the Varian V60 turbopump for use with our external ion source, and we are currently using that to pump the gas handler. It can bring the system from a few Torr to less than 0.1 mT in about 15 seconds (unless we have been using HD for a while, in which case it is necessary to flush the system with N_2 a few times to reach that level).

We have replaced the copper plumbing used for connecting additional lecture bottles to the system, because it was not very clean. We now have two lecture-bottle hookups available (currently holding HD and ^{3}He), and the tritium is loaded directly into a stainless steel sample bottle which replaces one of the sample bottles normally used (see Figure 2-9).

K. External ion source

The ultimate in fast ion-switching is to make the ions externally. Then, if we use a mass filter to select the ions we want, there will be no bad ions and no neutral gas load on the system. The ion source has been built, and produces a beam with enough mass resolution to separate out any bad ions we may make. Ideally, putting it to use will involve simply bolting it onto the experiment (actually it is fairly heavy and needs to be hung

from the ceiling as well) and tuning the steering electrodes until the ions wind up in the trap. Of course, things never go so smoothly, but there should be no serious obstacles in the way of making the ions externally. However, that will be the subject of somebody else's thesis.

III. HOW WE MAKE A MEASUREMENT

Our precision mass comparisons are still done by the "pulse-and-phase" method described in [CWB89] and [COR90]. This method uses a π-pulse to coherently swap the cyclotron amplitude and phase into the axial amplitude and phase, enabling us to read out the phase of the cyclotron motion. Thus if we pulse the cyclotron motion at t=0 and then π-pulse and read out the accumulated phase at $t = T$, we know $\varphi(T)$. Then we can do it again with a different T to find $\frac{d\varphi_c}{dT} = \omega_c'$ and to ensure that we unwrap the phase correctly. For more detail on this method, see Section III.B.2.

The dominant source of scatter in our cyclotron frequency measurements is (still) magnetic field fluctuations. Therefore, we (still) need to make our precision measurements early in the morning (1:00 to 6:00 AM), while the subway (~150 meters away) is not operating. We also have to disable the freight elevator down the hall, which causes more than 1mG of shift as it moves from the bottom to top floors. (One milligauss, divided by our shielding factor of ~8, is $10^{-8}B_0$.)

For more information about our magnetic field fluctuations, see [GAB90] and [CWB90b]. When we made the shielding measurements reported in [CWB90b], we found a much higher shielding factor (30±10) for the elevator than for other sources of noise. We believe this was because the magnetometer sensor was near the (steel) wall of our laboratory, and hence saw a larger field from the elevator than did the magnet. We have since moved the sensor away from the wall, and the measured field change as the elevator moves has been reduced by a factor of about 2. We now measure approximately the same shielding for all sources of magnetic field noise.

Figure 3-1 shows the shielding of our magnet with respect to the field at the (fixed) magnetometer station. We made a series of simultaneous measurements of external magnetic field and phase of an ion. The ion was atomic nitrogen (N^+), with a trap cyclotron

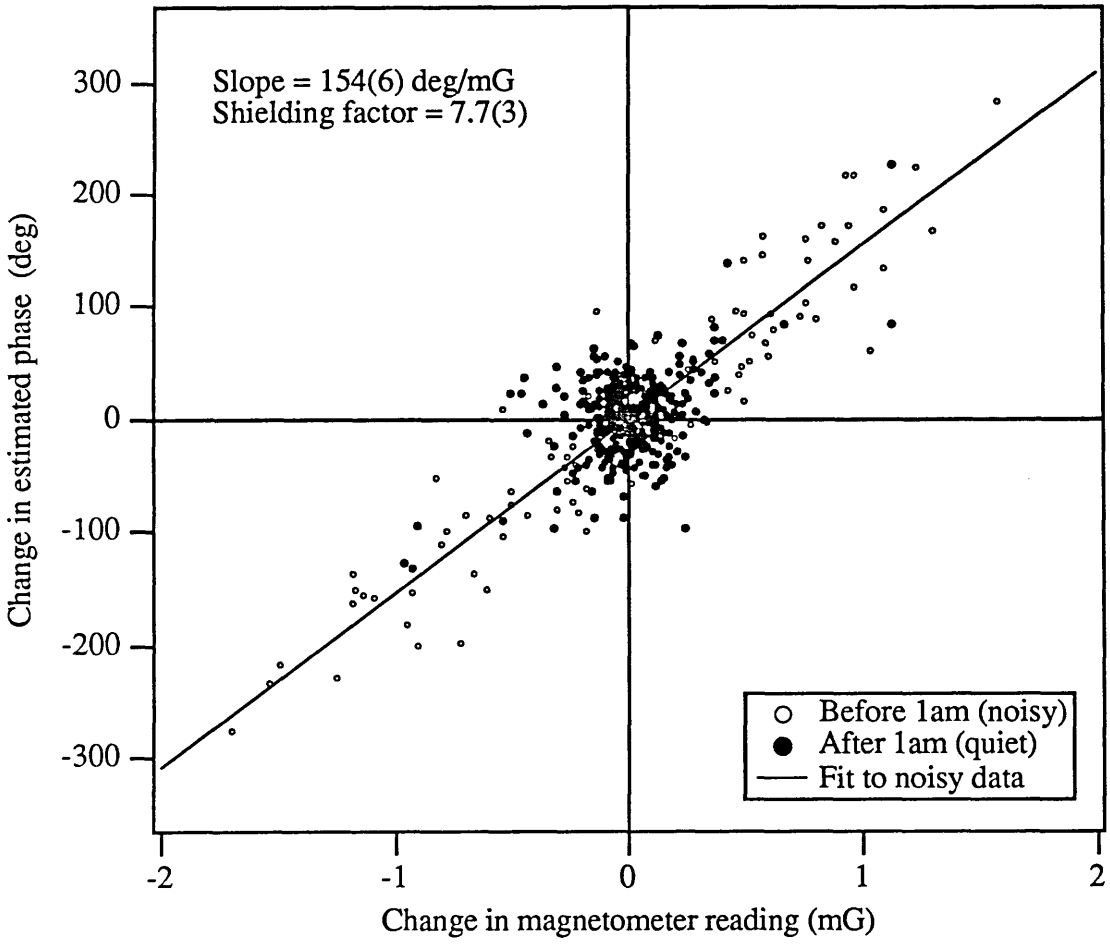

Figure 3-1. Magnetic shielding in the trap. The phase data is from a single N^+ ion, integrated for 30.200 seconds. The magnetometer is in the center of the magnet room, approximately 2m from the field center.

frequency of $\omega_c' = 9351570$ Hz, and the cyclotron frequency was integrated for 30.200 seconds.

The magnetometer reading drifts slowly (with a time constant of around 30 minutes), so we have to remove that drift before comparing to the ion frequency. Then, since we also remove any real drifts in the field, we need to remove long-term drift from the ion signal as well. We have done this in Figure 3-1 by using the differences between adjacent points. That is, the horizontal axis is the difference in measured B between two points, and the vertical axis is the difference in ion phase. We could instead fit the data

versus time to a polynomial (separately for the phase data and field data), and compare the residuals. We did this, and got the same shielding factor. The "quiet" points mostly cluster near the origin, so that no particular slope can be extracted. Note, however, that some of them are fairly far out, and line up well with the slope from the "noisy" time. These points correspond to times when the elevator was moved, whereas the noisy fluctuations are mainly due to the subway. Thus we know that the shielding factor is now the same for the elevator and the subway.

Since we have $\omega_c' = 9351570$ Hz, and the integration time was $T = 30.200$ s, one degree of phase shift corresponds to a fractional frequency shift of $1/(360T\omega_c') = 9.836 \times 10^{-12}$, which is equal to a field change of 8.385×10^{-4} mG (since $B_0 = 85250$ G). So the 154(6) degrees/mG that we measure corresponds to a field penetration factor of 0.129(5), or its inverse, a shielding factor of 7.7(3).

A. Preparation

Before we make a precision measurement, we need to "tune" the trap (that is, adjust the guard ring voltage to minimize C_4) and find the cyclotron and magnetron frequencies. The free-space cyclotron frequency is generally quite stable, of course, but the detector frequency and hence the axial frequency drift a couple of Hertz from day to day. Rather than calculate the trap cyclotron frequency $\omega_c'$and coupling frequency $\omega_d \equiv \omega_c' - \omega_z$ from a fixed ω_c, we measure ω_d by the avoided crossing method (see below and [CWB90a]). This ensures that we don't make an arithmetic mistake and that everything is properly connected to drive the cyclotron motion. Also, the magnetic field does change by about a part in 10^6 during the 6 weeks between fills of its liquid helium reservoir, which works out to about 10^{-9} per hour.

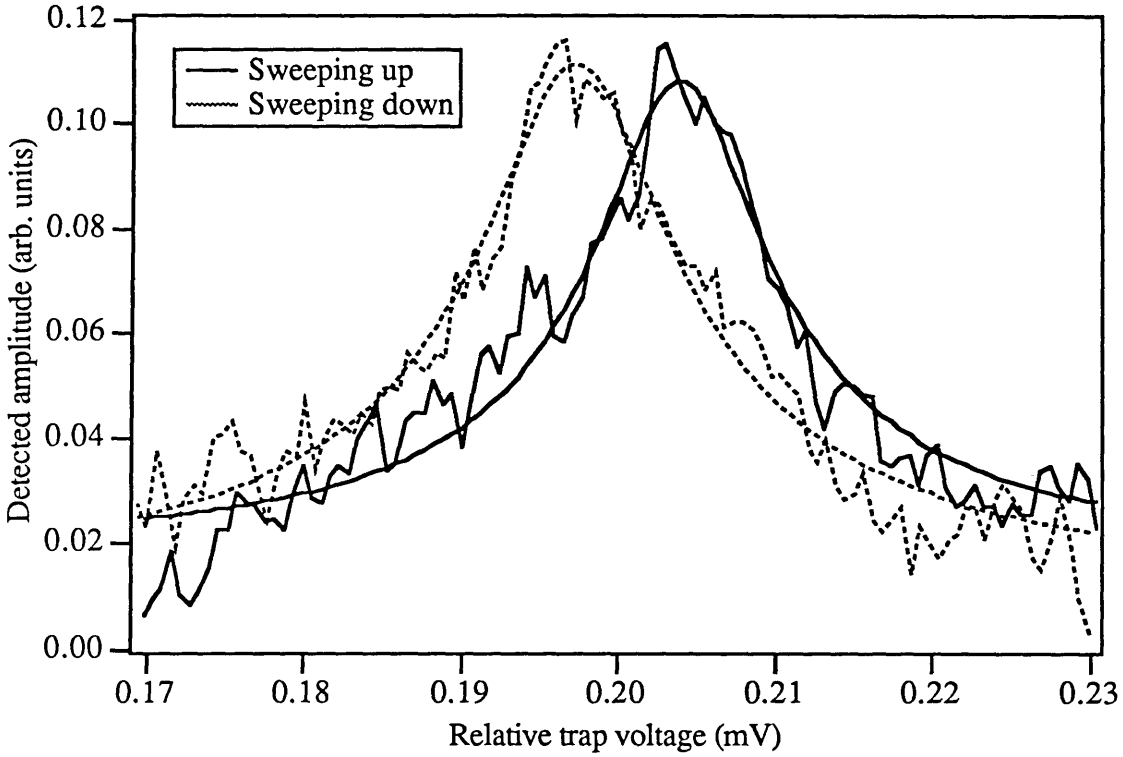

Figure 3-2. The ion resonance in a well-tuned trap, with a Lorentzian fit to each sweep. A 1 mV change of the trap voltage corresponds to roughly 9 Hz in frequency. The two sweeps don't line up because of the time constant of our lock-in amplifier.

1. Trap tuning

Before each daily measurement, we must carefully "tune" the trap. That is, we adjust the guard ring voltage, V_{gr}, to give the most harmonic ion response. This is to reduce C_4, which changes slightly from day to day, presumably as the charge patches change. Actually, as explained in [COR90], the most harmonic response does not occur at the minimum of C_4, but at the minimum of a linear combination of C_3 and C_4. However, since we have no way of measuring C_3, we have to assume it is small and attempt to minimize the anharmonicity.

Tuning the trap takes only about half an hour, and can be done before the ambient magnetic field settles down for the night, so there is no penalty in having to tune the trap

for each run. A typical trace which we consider "tuned" is shown in Figure 3-2. The smooth lines are Lorentzian responses, fit to the data. An anharmonic response would be asymmetric, "leaning" to one side or the other, or even showing hysteresis with the direction of sweep [LAL76, p. 89].

We took this data by the two-drive CW method detailed in [WEI88], which is essentially a lock-in scheme. We sweep the ring voltage in both directions, and adjust the guard ring voltage for maximum symmetry. Each sweep is offset (in the direction of the sweep) by the time constant of the lock-in amplifier, which is why the traces don't line up. This particular data corresponds to $C_4 < 2 \times 10^{-5}$, if we assume C_3 is less than about 3×10^{-3}.

2. Finding ω_c' by avoided crossing

We can find the cyclotron frequency to about 0.1 Hz by the avoided crossing method. This method was first described in [WEI88], and developed in the dressed-atom formalism in [CWB90a]. The basic idea is that the when the coupling drive is on, the two harmonic oscillator modes (cyclotron and axial) couple in such a way as to repel each other. That is, one has to find new normal modes of the system, which are a linear combination of the cyclotron and axial modes, and whose frequencies "anticross" as the coupling frequency is varied across the resonance. The frequencies of the normal modes as a function of coupling frequency are shown in Figure 3-3.

An example of an avoided crossing is shown in Figure 3-4, in the "waterfall plot" style commonly used to show atomic energy level anticrossings. This is the graphical complement to [CWB90a], which presents our (purely classical) avoided crossing in the language of atomic physics. Note that the peaks get smaller as they move away from the natural axial frequency. There are two reasons for this: the detector sensitivity is falling off, and the physical motion of the mode is becoming more radial and less axial. Since we detect only axial motion, we can only see the modes which have a reasonably large

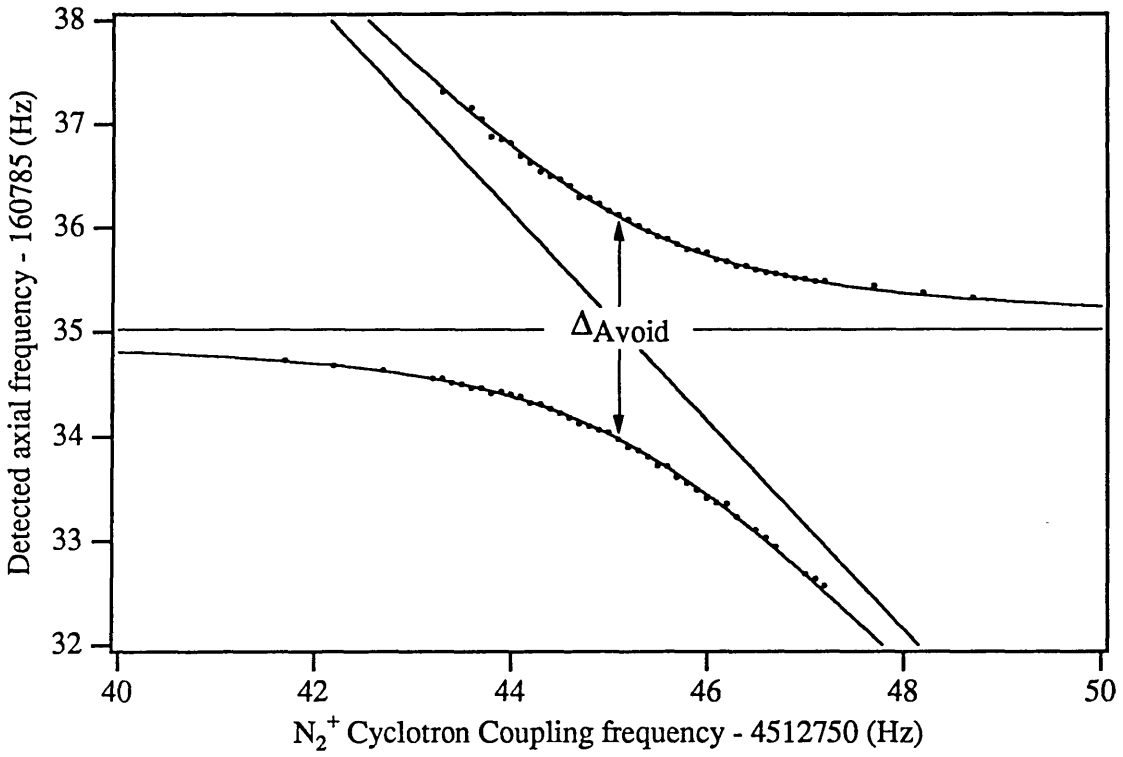

Figure 3-3. The frequencies of the normal modes versus coupling frequency. The asymptotes have slopes of 0 and -1. The points are data from Figure 3-4.

axial component. Each scan is a power (rather than amplitude) spectrum, since that makes the peaks stand out much better against the background noise.

We use the avoided crossing for three purposes: to find the cyclotron frequency when we don't know the strength of the magnetic field accurately, to check it before each measurement, as described above, and to find the amplitude-time product of the coupling drive which yields a π-pulse. I will describe each use below.

When the magnet quenches, as it has twice so far (once due to liquid nitrogen spilling onto the vacuum seal O-rings, and once, as mentioned above, necessitated by bringing a power supply too close to the magnet), we must recharge it. This results in some small uncertainty in the field strength. The first time we quenched the magnet, it was recharged by factory-authorized personnel with official equipment, but not carefully checked with an NMR probe, so there may have been a change of a few tenths of a

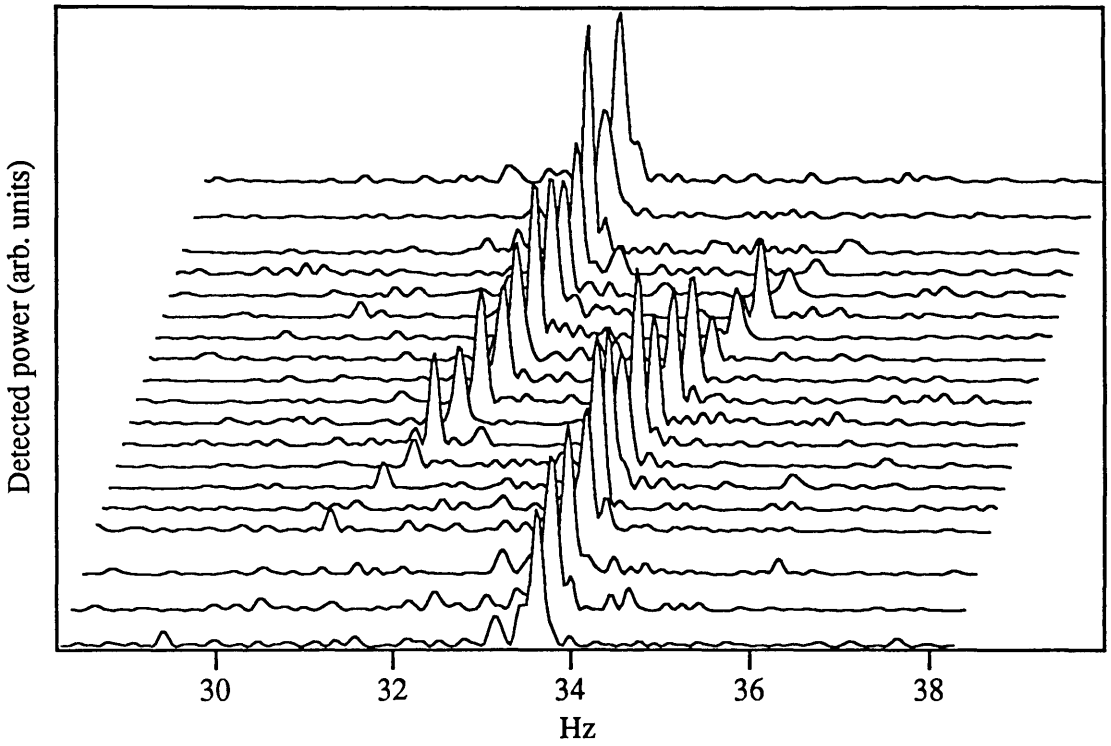

Figure 3-4. Each horizontal sweep is the power spectrum of the ringdown of one N_2^+ ion, pulsed to about 20% of the trap size. The horizontal axis is frequency, vertical is amplitude, and depth is frequency of the coupling drive. The spacing between the scans is not uniform because we take more scans closer to the resonance.

percent (we hadn't ever measured a cyclotron resonance before the quench). The more recent time, however, we recharged it ourselves, leading to an uncertainty of about 3 or 4% in the field. Therefore, we had to search over a wide range to find the cyclotron frequency. This is easy with the avoided crossing; the maximum coupling drive we can conveniently apply causes a detectable shift in the axial frequency even 100 Hz from resonance. Thus we just move the coupling frequency in 200 Hz steps and look for a shift in the axial frequency when we pulse the ion. This method can cover 1 kHz in a minute.

The third use for the avoided crossing is to calibrate our drive strength, and thus determine how many Volt-seconds to apply for a π-pulse. From [CWB90a] Equation 4, we see that, on resonance ($\delta = 0$), the frequencies of the axial modes are $\varepsilon = \pm\frac{1}{2}|V|$, where $|V|$

is the coupling strength, with units of frequency. That is, the splitting between the two frequencies on resonance is $|V|$. But we also know that the π-pulse criterion is $|V|\tau = \pi$, so if we use the same amplitude for the π-pulse as we did to measure the avoided crossing, we need $\tau = \frac{\pi}{|V|} = \frac{1}{2\Delta f}$, where Δf is the splitting in Hz. For the example shown in Figure 3-4, the splitting is 2.12(5)Hz, so we need τ = 236(6) mS.

We could of course map out the π-pulse directly by varying the coupling drive strength $|V|$ and plotting the amount of cyclotron action which gets transferred into the axial mode. This was done in [CWB90a], but it is not nearly so accurate as using the avoided crossing. This is because it uses the amplitude of the ion, which is harder to estimate than frequency. In addition, away from the correct π-pulse the amplitude is reduced (at 2π it's zero), further reducing our S/N. Thus we no longer map out the π-pulse directly as a function of coupling drive strength V. Instead, we do a careful avoided crossing and adjust the time as shown above. We do like to keep the same amplitude, in case there are nonlinearities in any of our attenuators or amplifiers. We have checked that our π-pulse transfers all the cyclotron action to the axial mode; indeed no detectable energy remains in the cyclotron mode even when the initial cyclotron radius is so large that the axial oscillation resulting from the π-pulse is shifted 2 Hz due to anharmonicity.

Measuring the strength of the coupling drive in this way is also important for ensuring that both ion species in a comparison are being driven with the same strength. If the transfer function of the electronics is not flat with frequency, we will see a different value of π-pulse for the two ions. Unfortunately, the two coupling frequencies are in general separated by something on the order of 1 kHz, whereas the cyclotron driving frequencies are 160 kHz away, so this check is imperfect. Nonetheless, it is how we discovered that we were initially operating on the steep edge of a resonance and indeed driving the ions to systematically different orbits. (See Sections II.D and V.B.2 for more information.)

B. The measurement cycle

After we have done the preparatory work and waited until the subway has shut down for the night (about 1:30 AM), we stop the elevator on the first floor and begin taking data. The data-taking procedure is entirely automated (for heavy ions, at least), and we are careful not to enter the magnet room during the measurement. We could almost let the experiment run by itself (indeed, we tried that once), but we need to be there to tweak some of the ion-making parameters, since the FEP current often changes dramatically over an hour. We also have provisions to pause the experiment, in case somebody needs to use the elevator. We return the elevator to the first floor before restarting.

The measurement consists of repeatedly making an ion and measuring its cyclotron and axial frequencies, alternating between the two types of ions we are comparing. Making the ion takes about 5 minutes, and measuring the frequencies takes less than 3 minutes. We actually make two precision (long integration-time) cyclotron measurements on each ion, at the beginning and end of the measurement time. This ensures that the time between "adjacent" measurements on alternate ions is minimized.

1. Making one cooled, killed ion

Making a single ion involves several steps. First we make a small number of ions (between 0 and 5), and thin them down to one, by "dipping" them close to the lower endcap with an antisymmetric axial voltage. Then we need to "kill" the ions of other species ("bad ions"), which we make whenever we make ions. Finally we cool the magnetron and cyclotron motions to bring the ion as close as possible to the center of the trap.

a. Making

The making of an ion is now totally automated, at least for mass 28 ions. Additional work will be needed to successfully make light ions such as $^3He^+$ or T^+

(which have a lower S/N ratio) or ions which require dissociation as well as ionization (such as T^+ or N^+).

The first step is to expel the old ion from the trap ("dumping"), which we do by briefly reversing the voltage on the trap. Next we make a batch of ions. This is still done in much the same way as was described in [COR90], but we have automated the process and added one very useful refinement—making the ions in a weak trap and adiabatically compressing them toward the center of the trap.

The basic idea is to make ions in the trap by bombarding neutral atoms with electrons. Since we now have computer control of the gas handling system and Field Emission Point, this process is controlled by the computer. All we have to do is push a button (actually click on a button icon), and the ions are made. This is important because it means we can make our entire measurement without entering the magnet room. Thus we have no chance of accidentally moving any of the various magnetic objects near the magnet while a mass comparison is under way.

The other improvement we have made is to make the ions in a very weak trap and then adiabatically increase the voltage, compressing the ions toward the center of the trap. Ions are presumably made uniformly along the z axis and so have a uniform distribution of initial axial amplitude ρ_z. Unfortunately, ions with a large ρ_z are so shifted in frequency (due to anharmonicity) that they don't couple to our detector and thus don't cool. Previously we simply smashed these ions into the lower endcap when we first dipped the ion cloud. We start with the trap voltage at about 10% of the voltage needed to bring the ions into resonance with the detector, then raise it to 100% over the course of 2 or 3 seconds. With this scheme, ions with any initial ρ_z are (to first order) compressed into the central 30% or so of the trap. This amplitude is small enough that the anharmonic frequency shift is less than the bandwidth of our detector, so the ions cool rapidly.

This procedure has made our ion-making much more efficient. We now need only about 5 mT-cc of gas with maybe 5 nA of electron current to make 2 or 3 N_2^+ ions. This

should be compared with roughly 50 mT-cc and 10 nA that we used with the old method. This is not so important for mass 28, which cryopumps perfectly, but for the mass 3 ions, vacuum contamination by neutrals was our bane in the past. Indeed, we have found that we can now consistently load single HD^+ ions, with no evidence of vacuum problems. Of course, helium may be worse than hydrogen in this respect, but we have certainly improved matters over our previous method.

After making a batch of ions, we have to reduce the number to one. This is done by dipping the ions close to the lower endcap. As mentioned in Weiskoff's thesis, we actually apply a voltage V_{LEC} only on the lower endcap. This can be decomposed into an antisymmetric voltage ($+V_{LEC}/2$ on the lower endcap and $-V_{LEC}/2$ on the upper endcap) and a symmetric voltage ($+V_{LEC}/2$ on both endcaps). The antisymmetric part shifts the ion cloud to a new equilibrium position ([WEI88], pp. 31-32)

$$\hat{z} = \frac{\sqrt{\omega_z^4 - \frac{3e^2 B_1 B_3}{m^2 z_o^4} V_{LEC}^2} - \omega_z^2}{\frac{3eB_3}{mz_o^3} V_{LEC}} \tag{3-1}$$

where e and m are the charge and mass respectively, B_1 and B_3 are geometrical constants, and z_o is the axial size of the trap, defined in Figure 1-1.

The symmetric part of the voltage on the endcap adds to the normal trapping potential to increase or decrease the strength of the trap. Thus if we apply a positive voltage, the trap gets deeper as the ions move away from the lower endcap, and if we apply a negative voltage, the trap gets shallower and the ions move toward the lower endcap.

In practice, we apply a fixed voltage to the lower endcap, so to vary the effective depth of the dip we change the ring voltage before we actually dip the ions. We can find a setting for the strength of the dip which will usually leave one ion in the trap. This setting is stable over the course of a week or so. Unfortunately, dipping to the canonical

voltage does not always leave us with one ion, so we need a way to count how many ions are in the trap.

b. Counting

The two parameters that characterize the ion pulse response are the width (damping) γ_i and the amplitude a. If the ion width is much less than the coil width, $\gamma_i << \gamma_c$ (our usual experimental regime), the ion damping is proportional to the number of ions. Thus we can use the width directly as an indicator of number of ions. If there are many ions in the trap, the pulsed response tends to become anharmonic, and the relationship between number of ions and width gets very noisy. Fortunately, we don't really care how many ions there are; we just want to know whether we have zero, one, or "more than one".

In addition to the width, the amplitude of the response grows slightly with the number of ions. Theoretically (see sections V.A and V.B of WEI88), the initial amplitude should go as

$$a = \frac{\gamma_i}{\gamma_o - 2\gamma_i} \tag{3-2}$$

which becomes, for $\gamma_i << \gamma_c$ (a condition certainly met for N_2^+),

$$\begin{aligned} a &\propto \gamma_i(N) \\ &\propto N \end{aligned} \tag{3-3}$$

so the initial amplitude of the signal is proportional to the number of ions N. However, when we count the ions by pulsing, we use a simple Fast Fourier Transform (FFT) to convert to the frequency domain, and we keep the sampling rate and number of samples constant. (Chapter 4 describes how we use the Laplace transform to convert to frequency domain when we need more accuracy.) We also adjust the sampling rate and number of samples to take data for approximately one single-ion damping time. Then when the damping doubles, half the samples are essentially just noise, and the normalized Fourier amplitude decreases. From WEI88 Equation V.B.1, we see that the peak amplitude of the

FFT goes as $1/\gamma_i$ (= $1/N$) times the peak amplitude a, when the data-taking time is much greater than the damping time.

Taken together, the considerations in the two previous paragraphs imply that the peak amplitude should be independent of the number of ions. However, we actually see a slight increase in signal going from one ion to two. This is presumably due to the fact that the FFT peak going as $1/\gamma_i$ is a fairly rough approximation in the regime of one or two ions, where the damping time is just greater than the data-taking time. That is, the contribution from $t = \tau_i$ to $t = 2\tau_i$ is only $1/e$ smaller than the contribution from the first damping time. Thus, increasing the data-taking time from τ to 2τ decreases the FFT peak height by 1.37/2, rather than 1/2. When combined with the factor of 2 increase in current, we should see 37% more signal for two ions than for one. This estimate is far more accurate than the signal from any single ion pulse, but the point is that the area under the frequency-space peak is more than twice as much (on average) for two ions as it is for one. Thus, even with the rather noisy signal we have to work with, we can identify when we have one (mass 28) ion with near 100% reliability (assuming the width-amplitude product for one ion has been properly calibrated).

For mass 3 ions, $\gamma_i \approx 5$ for one ion, and since with our Q of 32000 we have $\gamma_c \approx 30$, the condition of ion width << coil width (the condition that the ion and detector are weakly coupled, if you prefer; again see WEI88) doesn't hold even for 2 ions in the trap. However, we should still be able to distinguish between one and >1. In fact, we can do so visually, but implementing an algorithm to do that automatically may require more intelligence than we have been able to give our computer so far.

c. Killing Bad Ions

When we make ions in the trap, we always make ions of other species ("bad ions"). We have not tried to identify what the bad ions are, for the most part. Based on how much the critical voltage of the FEP changes from day to day, we strongly suspect that

tungsten ions are being made, but we would need about 60V on the trap to see W^+ (atomic weights of 182-186). We also know that for diatomic molecules like N_2 and CO, we make some dissociation products. The situation is substantially worse when we try to make atomic ions by dissociating molecules, e.g. N^+ from N_2 or T^+ from T_2. Then the undesired molecular ions outnumber the atomic ions, by about 10 to 1 in the case of nitrogen.

Whatever the source of bad ions, and whatever their species, we need to eliminate them from the trap. We do this by driving the bad ions (but not the one we wish to keep) with shaped white noise, and then dipping the entire cloud close to the lower endcap [COR90]. Previously, we made the white noise with our computer, which limited us to a bandwidth of about 60 kHz. Thus we had to use the noise to modulate a carrier, and step the carrier frequency until we had covered all the frequency space up to whatever corresponds to mass 1 (about 850 kHz when the trap voltage is set for N_2^+). We now use a wideband noise generator and notch filter (see Section II.E.2), which reduces the killing time by a factor of 15 to 20. This is a totally straightforward change to the apparatus, yet it is by far the most important factor in increasing the number of measurements per night from 3 (just enough to be useful) to 40 (enough to do real statistics on).

We have also developed a slightly different technique for getting rid of the large number of molecular ions which are created when we make a monatomic ion from diatomic neutral gas. In this case, we know quite accurately what frequency the bad ions will have, so we can pump lots of energy into them without disturbing the good ion. We still need to spread the energy slightly in frequency space, because the anharmonicity would otherwise limit the maximum orbit size of the bad ions. So we apply a carrier at the appropriate frequency ($f_o/\sqrt{2}$, where f_0 is the frequency of the good ion), modulated by noise with a bandwidth on the order of 1 kHz.

We find that if we just try to kill the normal way, we either lose the good ion or can't get rid of all the bad ions. This is probably because of the very large number of bad

ions, and the resulting stronger coupling to the good ion. If we get all the bad ions of a particular mass (e.g. the molecular ions) moving in phase, they will look like a single very massive, highly charged ion and transfer some of their energy to the good ion before we dip them. If, as seems more likely, the bad ions are not all in phase, we will have to apply more drive to make sure all of them are in large enough orbits to be neutralized when we dip the cloud. In that case, the additional wideband noise will excite the good ion directly (a deeper notch filter would help).

In either case, driving only the bad ions and immediately dipping them reduces the direct excitation of the good ion to an insignificant amount, and prevents the energy in the bad ion cloud from transferring to the good ion. Then, when we dip, the difference in orbit sizes between good and bad ions is as large as possible, causing the bad ions to leave while preserving the good ion.

Of course, in this case we can check directly whether there are any bad ions remaining, since we know the frequency exactly. This is in fact how we learned that the standard killing procedure was not effective, and that direct excitation of the known bad ions was necessary to eliminate them completely. We also do the standard killing procedure afterwards, to remove the usual suspects (tungsten, etc.).

d. Cooling

Having made a single ion of the desired species, and ejected all others, we need to cool the radial motions of the ion. Cooling the cyclotron frequency is easy, since we know the frequency and time-amplitude product for a π-pulse already (see Section III.A.2). With the ion frequency tuned to the detector resonance, we wait 3 or 4 damping times for the axial motion to die out, then apply a π-pulse to swap whatever cyclotron action the ion might have into the axial mode, and wait for that to die down. Actually, we do this right after the final dip of the killing routine, so the ion is already cooled axially.

Thus the whole procedure takes only about 4 damping times, or about 15 seconds for mass 28.

In principle, we could do the same thing for the magnetron motion, but in practice we use simple "sideband cooling" [VSD78, WID74, WID75]. This means simply leaving the coupling drive on continuously, letting the action slosh between axial and magnetron modes and dissipate in the detector due to the axial component of the motion. We do this because the maximum amplitude of magnetron coupling we can apply is very small, due to the design of our cryofilters. Thus the maximum splitting Δ_{avoid} of an avoided crossing is small compared to the ion width, and hence hard to measure accurately.

2. Measuring the cyclotron frequency

The basic idea for measuring the cyclotron frequency ω_c' is the same "pulse-and-phase" method described in [CWB89] and [COR90] (and briefly below), but we have added a technical refinement to improve our frequency precision. This, combined with our increased S/N and our improved procedure for extracting the phase from the detected signal (see Chapter 4) has reduced the time required for a measurement dramatically. For an individual phase measurement of given precision, the time needed has been reduced by a factor of about 4. The overall measurement time has been further reduced, because we no longer need to reject 20% of our phase points as too noisy. For the mass 28 measurement at a precision of 2×10^{-10} per shot (and $T = 40.2$ sec), these two considerations and the shorter ion-making time combine to reduce the measurement time by over a factor of 10.

The pulse-and-phase method is quite simple in principle: drive the cyclotron motion with a short pulse of frequency $\omega_d \approx \omega_c'$ at $t = 0$, let the ion's cyclotron phase evolve at the rate $\varphi(t) = \omega_c' t + \varphi_0$ for T seconds, then π-pulse it into the axial motion and read out the phase. Then the average cyclotron frequency is $\omega_c' = \frac{\varphi(T) - \varphi_0}{T}$, which is fine except we don't know φ_0. So repeat with a different T, and we have the cyclotron fre-

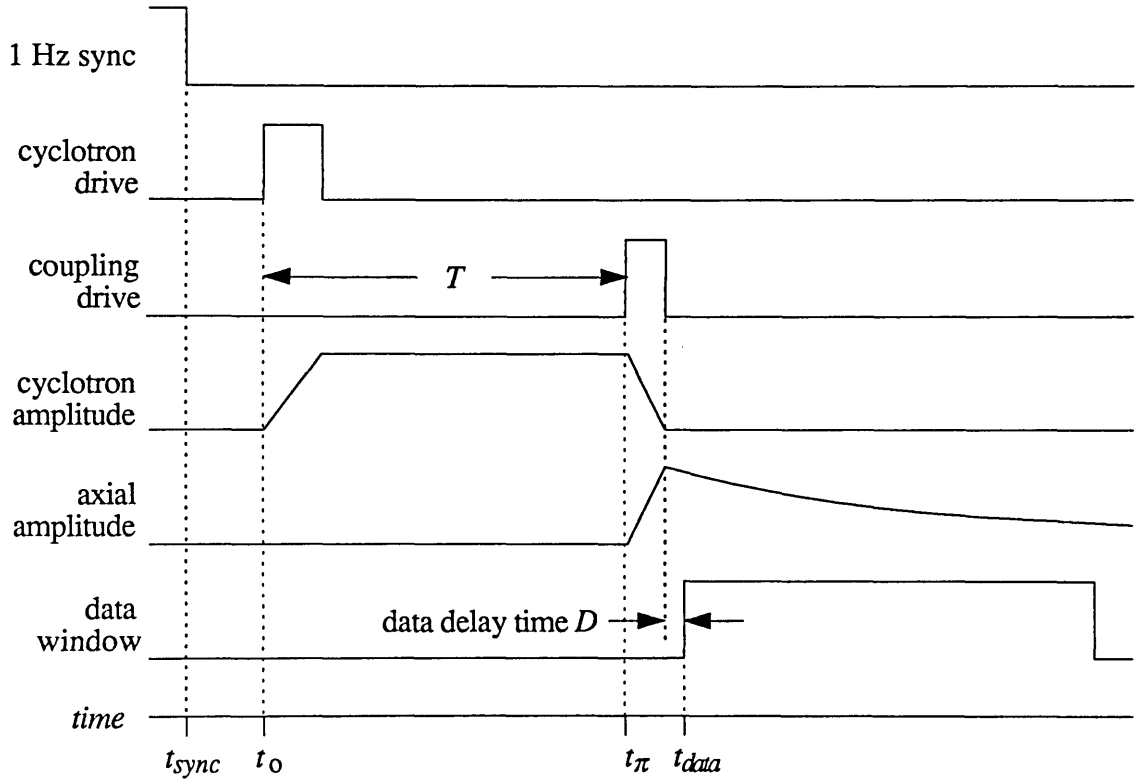

Figure 3-5. The timing of a pulse-and-phase measurement (not to scale). T is the phase integration time. D is typically 2 mS.

quency $\omega_c' = \frac{\Delta\varphi}{\Delta T}$. Naturally we need to be careful to unwrap the phase correctly, but this can be done unambiguously by taking a suitable number of data points at intermediate times. Figure 3-5 shows the timing of the important drives and amplitudes.

Once the phase has been unwrapped, we find the frequency by fitting a line to $\varphi(T)$. The points at the ends of the line have the largest lever arm, so to first approximation the uncertainty in frequency is

$$\delta\omega \approx \frac{1}{T_{max} - T_{min}} \sqrt{\delta\varphi^2(T_{min}) + \delta\varphi^2(T_{max})} \qquad (3\text{-}4)$$

where T_{max} and T_{min} are the longest and shortest phase integration times, respectively. In reality, each successive T is about 3.5 times longer than the previous one, so while this approximation is good at the T_{max} end, two or three of the shortest Ts contribute to the fit

at the low end, Thus we have to use $\delta\varphi_{eff}(T_{min})$, which is less than $\delta\varphi(T_{min})$ due to the additional data points near T_{min}.

We can, in fact, reduce the error by measuring φ_0 directly. To explain this we need to delve into the details of the (many) phases and how they relate to each other. We see from [COR90] that, after the π-pulse, the complex amplitude of the axial motion is given by:

$$Z = \frac{P^*}{|P|} C_0 \tag{3-5}$$

where C_0 was the complex amplitude of the cyclotron motion before the π-pulse ($t = t_\pi$ in Figure 3-5) and P is the complex amplitude of the coupling drive. If we use polar coordinates, so that, for example, $Z = Ze^{i\varphi_z}$, the phase portion of Eq. (3-5) becomes

$$\varphi_z = \varphi_c - \varphi_p. \tag{3-6}$$

That is, the phase we detect is actually the phase difference between the cyclotron motion and the coupling drive, at the time of the π-pulse.

Finally, we need to consider the phase of the mixer (see Figure 2-6.). We multiply all our signals by the mix frequency ω_{LO} ("LO" stands for "Local oscillator"), to bring them down to a reasonable frequency for data-taking (typically around 35 Hz). The effect of this on the detected phase is that we have to add back the phase of the mixer to get the phase of the ion signal. To see this, consider the effect of the mixer. We multiply the signal ω_z and local oscillator ω_{LO} to get the "intermediate frequency"[6] signal:

$$y_{IF} = \cos(\omega_z t + \varphi_z)\cos(\omega_{LO} t + \varphi_{LO}). \tag{3-7}$$

Recalling that

$$\cos\alpha\cos\beta = \tfrac{1}{2}\cos(\alpha-\beta) + \tfrac{1}{2}\cos(\alpha+\beta), \tag{3-8}$$

---

[6] A misnomer, but I'm running out of subscript mnemonics.

we find, after filtering out the high-frequency term, that the signal at the input to the A/D converter is

$$\begin{aligned} y_a &= \tfrac{1}{2}\cos\left[(\omega_z t + \varphi_z) - (\omega_{LO} t + \varphi_{LO})\right] \\ &= \tfrac{1}{2}\cos\left[(\omega_z - \omega_{LO})t + \varphi_z - \varphi_{LO}\right]. \end{aligned} \tag{3-9}$$

The initial phase of y_a at the start of data-taking is thus

$$\begin{aligned} \varphi_a &= (\omega_z - \omega_{LO})D + \varphi_c - \varphi_p - \varphi_{LO} \\ &= (\omega_z - \omega_{LO})D + \omega_c T + \varphi_0 - \varphi_p - \varphi_{LO} \end{aligned} \tag{3-10}$$

where D is the data delay time, and T is the phase accumulation time, defined in Figure 3-5. We can ignore any terms which are constant, since we always look at the slope of phase versus time. For instance, we ignore the first term, since a 1 Hz shift in f_z (which would cause no end of other problems) times our standard D of .002 second results in less than 1° of phase shift. Then, since the term we are interested in is $\omega_c T$, we need to find only the last three terms.

Now, since the initial phase of the cyclotron motion is determined by the phase of the pulsing drive ω_d, we can measure φ_0 very simply. At $t = t_0$, the ion is in phase with the pulsing drive. This implies that $\varphi_0 = \varphi_d$, but we need to be careful how we define the phase of the driving signal. φ_d is the phase of the driving *field*, which is not the same as the phase of the signal at the output of the driving signal generator, since there are various unknown phase shifts in the cryoelectronics. Let us write the phase of the driving field as

$$\begin{aligned} \varphi_d &= \varphi_{ds} + \varphi_{dx}, \\ \varphi_p &= \varphi_{ps} + \varphi_{px}, \end{aligned} \tag{3-11}$$

where φ_{ds} is the phase of the driving voltage at the signal generator (which we can measure, as shown in Figure 3-5), φ_{dx} is the sum of all the unknown phase shifts in the cyclotron driving electronics, and φ_{ps} and φ_{px} are similarly defined for the coupling drive. What we can measure directly is the phase relationship between our three signal generators at $t = t_0$,

$$\begin{aligned}\varphi_{HP} &\equiv \varphi_{ds} - \varphi_{ps} - \varphi_{LO} \\ &= \varphi_d - \varphi_{dx} - \varphi_p + \varphi_{px} - \varphi_{LO} \\ &= \varphi_d - \varphi_p - \varphi_{LO} + \varphi_x ,\end{aligned} \tag{3-12}$$

where $\varphi_x = \varphi_{px} - \varphi_{dx}$ is the sum of the electronic phase shifts. The first three terms are what we want to know, so we need only to find φ_x. Since measuring it directly (especially at cryogenic temperatures) is virtually impossible, the obvious solution is to make many pulse-and-phase measurements with two very small values of T, and subtract the extrapolated value of φ_0 from φ_{HP} to get φ_x. Then, as long as we don't change cables or otherwise disturb the phase shifts in the driving and detecting systems, we can then add φ_x to φ_{HP} to determine φ_0.

Since we still have to measure φ_0 using the ion, it might seem that we don't buy anything in the ultimate uncertainty of ω_c'. However, we do win because we can now spend 15 minutes or so measuring φ_0 very carefully at the beginning of a mass comparison (even before the magnetic field has become quiet), and then before every measurement we have only to measure φ_{HP}, which takes under 10 seconds. We wouldn't even need to re-measure φ_{HP} if we could keep the oscillators running throughout the run, but then we would need four oscillators rather than two. In fact, we also use the same signal generators for the magnetron cooling after making an ion, so we would need 6 signal generators instead of the two we have.

One point that remains is: how is $t = t_0$ defined? Well, since the three frequencies ω_d, ω_p, and ω_{LO} are constant, t_0 can be any time when φ_{HP} has some given value. To make sure this is true, we start by slaving all the clock circuits in the experiment to a single master clock[7], which is specified as stable to better a part in 10^{12} over 100 seconds, and 5×10^{-11} per day. Then we use only integer frequencies for ω_d and ω_p, and synchronize all our data-taking with a 1 Hz clock (see Figure 3-5 for timing, and WEI88, Section

[7]FTS 1002A, from Frequency and Time Systems, Inc., 34 Tozer Rd, Beverly, MA 01915

IV.E for a description of the synchronizer, although we have simplified the pulse synchronizer circuit slightly). Then ω_d, ω_p and ω_{LO} will always have the same relative phase at $t = t_0$, and it is this phase which we measure.

C. Converting to neutral atom values

After taking the data, we must go through several steps of analysis to find the mass ratio we are measuring. This consists of: (1) extracting the ion's frequency and phase; (2) unwrapping the cyclotron phase; (3) converting the trap frequencies to free-space values; (4) correcting for temporal drift of the magnetic field; and (5) taking the ratios. Parts (2)-(5) are straightforward, and described by example in Chapter 6, while part (1) is described in Chapter 4. Actually, however, there is one more step; converting the ratio of ion masses to the ratio (or difference) of neutral atom masses. I will describe that here.

We always measure the masses of ions, while the values of interest are in general the masses of neutral atoms. It is of course straightforward to convert the ion mass ratios to neutral-atom values. However, in light of the fact that we have done the conversion incorrectly in the past (most notably in [CWB90b], where we did not properly include the ionization energy), it is worth mentioning the procedure here. Mass doublets are typically reported in the literature as differences, so I will describe that conversion.

Consider two species X_1 and X_2. We define the masses (in atomic mass units) of the atoms and ions as

$$\begin{aligned} M_i &\equiv M(X_i), \\ M_i^+ &\equiv M(X_i^+). \end{aligned} \tag{3-13}$$

We measure the ratio of the ion masses

$$R \equiv \frac{M_1^+}{M_2^+}, \tag{3-14}$$

whereas in general, the number of physical interest is the mass difference between the neutral species,

$$D \equiv M_2 - M_1. \tag{3-15}$$

We must account for the mass and ionization energy of the electron, using

$$\mathrm{X}_i + E_i \rightarrow \mathrm{X}_i^+ + e^-, \tag{3-16}$$

where E_i is the first ionization energy (a positive number). The mass relation from (3-16) is

$$M_i = M_i^+ + M_e - E_i, \tag{3-17}$$

where E_i is in units of amu/c^2. We can now write the mass difference as

$$\begin{aligned} D &= M_2^+ + M_e + E_2 - M_1^+ - M_e + E_1 \\ &= M_2^+ - M_1^+ - E_2 + E_1 \\ &= M_2^+(1-R) - E_2 + E_1 \\ &= (M_2 - M_e + E_2)(1-R) - E_2 + E_1 \\ &= (M_2 - M_e)(1-R) - RE_2 + E_1 \end{aligned} \tag{3-18}$$

where we have used (3-14) and (3-17).

The error in D will be dominated by the error in either M_2 or (1-R). For a mass doublet, the value of (1-R) is typically a few $\times$ 10^{-4}, so the accuracy of M_2 will not be a consideration unless it is a few thousand times worse than the accuracy of R. As an example, the mass of N_2 is known to 2 parts in 10^9, so it will not limit us at accuracies above 1×10^{-12}.

IV. NEW DIGITAL SIGNAL PROCESSING ALGORITHM

One of the major improvements in the experiment has been the introduction of a new algorithm for extracting the ion's frequency, phase, and amplitude from our data. We discovered that the method we had been using was not only noisy but biased. That is, in the limit of many measurements, the average estimated frequency did not always converge to the correct value.

Vasant Natarajan found that the method of maximum likelihood is unbiased and approaches the theoretical ("Cramer-Rao") limits for uncertainty. The basic idea is a fairly standard technique in digital signal processing (DSP), and adapts nicely to our situation. We have implemented this algorithm for all our data analysis. In this chapter I will demonstrate the bias in our old method, describe the new method, and show evidence for the improved performance of maximum likelihood over our old FFT-based algorithm.

A. What we need to extract

When we are operating in a pulsed mode (which we do except while tuning the trap), there are four parameters we wish to estimate from the data. We model the ion's response as a damped sinusoid

$$y(t) = Ae^{((i\omega-\gamma)t+i\varphi)} \tag{4-1}$$

which is defined by its frequency ω, initial amplitude A, initial phase φ, and damping rate γ. This assumes that the ion is weakly coupled to the detector (see Section III.B.1.b), which is true for mass 28 or 18, but marginal for mass 3. It further assumes that the anharmonic frequency shift is negligible, a condition which is not met. We have seen the effect of frequency shift with amplitude in our precision mass comparisons, and this is discussed in Section VI.E.

In practice, γ does not vary much over the course of a day, since it is given by

$$\gamma = \frac{1}{m}\left(\frac{eB_1}{2z_0}\right)^2 \mathrm{Re}(Z_{det}), \tag{4-2}$$

where m is the ion mass, e is its charge, B_1 and z_0 are trap parameters, and Z_{det} is the impedance of the detector. On resonance, $Z_{det} = \omega_z LQ$, where ω_z is the resonant frequency, L is the detector inductance, and Q is the quality factor of the detector resonance. Since ω_z and L are constant, γ changes only if the Q changes, and once we have measured the damping for a given Q, we can easily find the damping at any time by measuring the Q. Thus we can consider γ to be known, leaving the three parameters ω, A and φ to be determined.

The first thing we do with the data is to sample it at the rate f_{samp}, taking it into the discrete-time domain:

$$y_n = Ae^{((i\omega-\gamma)nh+i\varphi)} \tag{4-3}$$

where h is the interval between samples, so $h = 1/f_{samp}$, and n runs from 0 to N - 1, where N is the number of samples. This brings up the usual problem of aliasing, where the frequency components from $f_{samp}/2$ to f_{samp} are folded back to the range $f_{samp}/2$ to 0, and each additional frequency band from nf_{samp} to $(n+1)f_{samp}$ also appears, in the same way, in the first $f_{samp}/2$ range.

To avoid aliasing, we use a four-pole Chebychev lowpass filter (see WEI88, Section IV.E), with a cutoff around 500 Hz. This means that when we take data at f_{samp} = 1000 Hz, the components in the first aliased band are not fully attenuated, and show up in the digitized signal. However, we always keep the signals of interest below 100 Hz, where the aliased components come from between 900 and 1100 Hz. The filter attenuation at 900 Hz is about a factor of 20, so there is a minimal amount of aliased signal (white noise, in general) added to our signal. Therefore, we ignore aliasing, as I will do in this chapter.

B. Straight FFT with debinning

The most straightforward method of spectral estimation is to perform a Fast Fourier Transform (FFT) and use the frequency, phase and amplitude of the peak bin directly. This leads to an extremely coarse estimate of frequency, and biased estimates of phase and amplitude. A complete treatment of FFTs, including the problems they cause for parameter estimation, is given in [OPS75] and [WEI88], so I will give just a brief overview here, to define some terms.

The FFT is just a fast implementation of the Digital Fourier Transform (DFT),

$$X_k = \sum_{n=0}^{N-1} x_n e^{-2\pi i n k/N}. \tag{4-4}$$

Time is indexed by n, and frequency is indexed by k. We refer to the individual X_k's as "bins". Notice that the spectrum X has no explicit scaling; the DFT merely takes an array of equally-spaced samples in the time domain into another array in the frequency domain, without regard to the sampling rate. To convert back to actual frequency we need the bin spacing, which is just the inverse of the data-taking time:

$$f(k) = \frac{k}{T_{data}}, \tag{4-5}$$

where

$$T_{data} = \frac{N}{f_{samp}} = Nh. \tag{4-6}$$

For instance, if we take data for one damping time (which provides roughly the optimal S/N for FFT analysis), the frequency resolution is the inverse of the damping time. (The damping time $\tau = 2/\gamma$ is the time for the *amplitude* to decay by a factor of $1/e$), So for a damping time τ of 4 seconds, we can determine the frequency to only 250 millihertz, which is a fractional accuracy of only 1.5×10^{-6}. This contributes an uncertainty of 2×10^{-9} for mass 28, so we clearly need at least an order of magnitude improvement.

The first solution, described by Eric Cornell in his thesis, is to combine information from the two highest amplitude bins of the FFT. The frequency is some fraction of the distance between the peak bin and the second highest bin, and the ratio of the two amplitudes determines that fraction. Using the terminology from [COR90], pp. 20-23, we define the ion's frequency in units of the FFT bin width as

$$f_{ion} = \frac{k_{ion}}{T_{data}} = \frac{k_0 + \varepsilon}{T_{data}}, \tag{4-7}$$

where k_0 is an integer, and $|\varepsilon| < 0.5$. We can then write the ratio of the amplitude of the two peak bins as

$$R \equiv \frac{A_0}{A_1} = \sqrt{\frac{(1+\varepsilon)^2 + \beta^2}{\varepsilon^2 + \beta^2}}, \tag{4-8}$$

where A_0 and A_1 are the amplitudes of the peak bin and second-highest bin respectively, and β is proportional to the ratio of the data-taking time T_{data} to τ:

$$\beta = \frac{T_{data}}{\pi\tau} = \frac{N\gamma}{2\pi f_{samp}}. \tag{4-9}$$

I have also ignored the sign of ε since everything is antisymmetric around 0.

As mentioned above, to maximize the S/N ratio for this method we take data for one damping time, so β will be about $1/\pi \approx 0.3$. The curve of ε vs. R for this condition is shown in Figure 4-1-a. Note that R is largest when ε is near 0. The problem is that the maximum value of R that we can measure is limited by noise. With our present detector, we have a signal-to-noise (S/N) of ~6; the rms noise amplitude in each bin at the peak of the detector resonance is about 1/6 the signal amplitude. (This is an unusual definition of S/N, but it is appropriate for our use. Our signal is a single damped sinusoid, so the noise bandwidth for our purposes is only one or two bins.) Thus, if the frequency happens to fall right on a bin ($\varepsilon \approx 0$), so that the correct value of R is 3.2, then the noise adding in

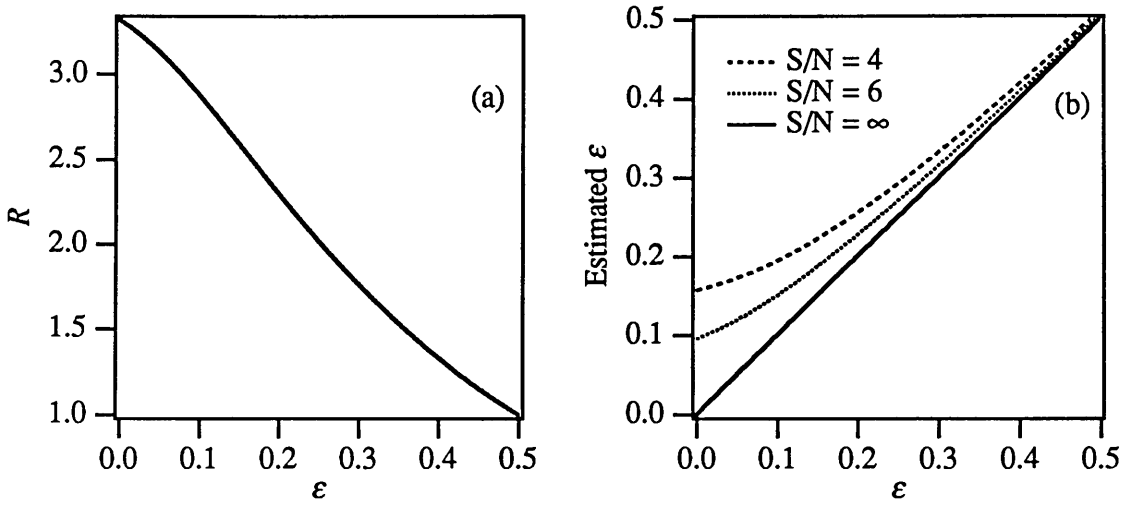

Figure 4-1. (a) R, the ratio of the amplitude in the two largest bins, vs. frequency. Frequency is in units of the bin width, f_{samp}/N. (b) The frequency estimate plotted against the true frequency, for no noise and for our current signal-to-noise ratio (6) and the Trap 2 S/N (4).

quadrature gives $R = 2.9$, which corresponds to $\varepsilon \approx 0.1$. That is then the lowest value that ε can take on using this estimation method.

This is what is meant by the term "biased estimator"; for frequencies which happen to fall on or near an FFT bin, the expectation value of the estimated frequency is not the true frequency. We can see this in Figure 4-1-b, where the estimated frequency is plotted against the true frequency for no noise, for a S/N of 6 (our current value), and for a S/N of 4 (the value we had with Trap 2).

If ε is negative (the second-highest bin is at a lower frequency than the peak bin, rather than at a higher frequency), the graphs in Figure 4-1 are reflected through the origin. That means the effect of noise is to map all frequencies with $|\varepsilon| < 0.1$ to $\varepsilon \approx \pm 0.1$, where the sign of the result could be positive or negative, though it is biased in the correct direction. Thus for 20% of possible frequencies, the uncertainty in frequency is 0.2 bins, which is 50 mHz for the canonical 4 seconds of data. This is still unacceptably large,

since it leads to a cyclotron frequency uncertainty of 4×10^{-10}, and because it is an artifact of data analysis, and not inherent in our data.

C. Zero padding

It turns out there are many methods in use by the Digital Signal Processing community to do spectral estimation [KAM81], [KUT82], and straightforward FFT is one of the worst. One conceptually simple, though calculationally inefficient, method is to pad the data with zeros before doing the FFT. This results in a spectrum with more points in the same frequency range, or "higher resolution". Of course no additional information is actually extracted, so the resolution is not truly increased.

Figure 4-2 shows the effect of zero-padding on a sine wave with noise (but, for clarity, without damping). Note that without padding, the second-highest bin is at a lower frequency that the peak bin, although the signal frequency (5.0 Hz) is actually higher than

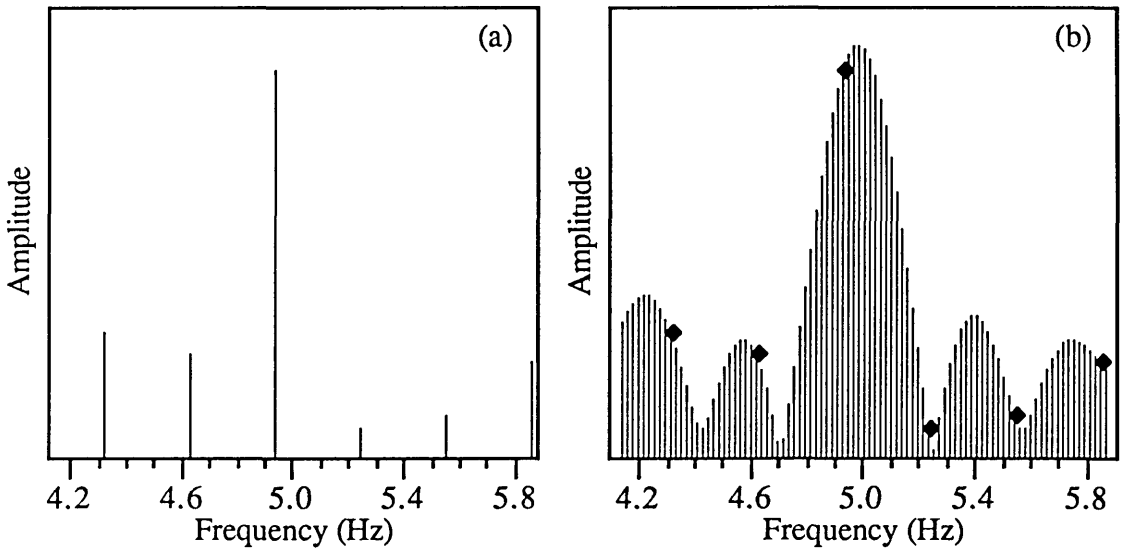

Figure 4-2. The effect of zero-padding on the amplitude of an FFT. The time-domain data is 64 samples of a sine wave at 5.0 Hz plus Gaussian white noise. (a) Amplitude of the 64-point FFT. (b) The same data after zero padding to 1024 points. Every 16th point corresponds to a bin in the 64-point FFT; these points are marked with diamonds.

the peak bin frequency. Thus the debinning procedure described above would yield a frequency estimate nearly half a bin low. Zero padding before doing the FFT correctly interpolates between the bins, giving a much better estimate of the frequency .

We use zero padding followed by FFT for the "real-time" display of data, which we use as a sanity check while the experiment is running. The use of zero padding improves our "eyeball" estimates of the frequency and amplitude substantially over straight FFT, making it easier to see when something is wrong. (Yes, I said zero padding was inefficient. We pad only to 2 or 4 times the length of the data. This yields a good graph for eyeballing, but we would need to pad by a factor of 16 or 32 (resulting in a 64K or 128K FFT) to get a precise value this way.) We save all data in straight time-domain form for off-line analysis as described below.

D. Maximum likelihood

We have chosen the method of maximum likelihood [KAM81] as a computationally reasonable and highly accurate algorithm to extract the parameters of the ion signal. This method was researched and implemented by Vasant Natarajan, who will describe the method in detail in his thesis.

The basic idea is to pick a model for the data, and then fit the parameters of the model to the data. We model the data as a damped sinusoid, and find the best estimate for ω, A, and φ by an iterative least-squares fit. The actual calculations are straightforward; we multiply the data by a decaying exponential (using the known decay rate γ), then find the frequency of maximum spectral amplitude by iteratively calculating the Discrete-Time Fourier Transform (DTFT). (The DTFT maps a discrete time-domain signal into a continuous frequency domain, unlike the DFT, which is discrete in both domains.)

The result of using this method is that our phase estimate now has a standard deviation of about 10°, roughly half of what it was before. The frequency estimation is

harder to characterize from actual data, because there is some actual scatter in the ion's frequency, due to fluctuations in the trap voltage. The standard deviation over short times (where drift of the trap voltage is negligible) is about 15 mHz for mass 28 ions. (Lighter ions have a lower S/N ratio, so the uncertainties will increase). This is an improvement over the old method, but most important is the elimination of bias in the estimation. This means we can safely use the estimated frequency, without having to add excessively large error bars to account for the biasing effect.

As an example, consider mass 28. The factor of two improvement in phase accuracy translates directly to a factor of two in ω_c'. Thus in a 40 second measurement we achieve a precision of $(10°/360°)/(40\,\omega_c') = 1.5 \times 10^{-10}$. The axial frequency adds (in quadrature) a similar uncertainty, since we know ω_z to only $.015/161000 = 9.4 \times 10^{-8}$, which must be multiplied by $(\omega_z/\omega_c)^2$ to get 1.1×10^{-10} before combining with the fractional error in ω_c'. Thus an individual 40 second measurement now has a precision of 2×10^{-10}.

V. SYSTEMATIC SOURCES OF ERROR

There are many sources of error, both random and systematic, in any measurement at this level of precision. Most of the errors in our experiment have been discussed in detail in [COR90], whence I refer the reader for further information. In this chapter I will examine the systematic errors, particularly those that were not considered previously, or whose parameters have changed in the process of improving the experiment. The statistical errors will be dealt with in Chapter 6, along with our actual data.

The sources of systematic error fall into two broad groups; amplitude-dependent and amplitude-independent. The former are proportional to some power (generally two) of one of the three amplitudes of ion motion in the trap, a_z, a_c, or a_m. The latter are not dependent on the amplitude of the ion's motion, and thus cannot be extrapolated to zero amplitude. I will first consider the amplitude-independent sources of error.

A. Amplitude-independent errors

There are four amplitude-independent sources of error that I will consider: pulling of the axial frequency by the detector; systematically different numbers of bad ions; shift in the estimated phase due to anharmonicity; and patch effects pushing the ions to different locations in the trap, where the magnetic fields are different.

1. Tuned circuit pulling

The tuned circuit in our detector interacts with the ion, causing damping of the ion and shifting its frequency. I refer the reader once again to [WEI88] for a complete treatment of the ion-detector system. Using the present detector we are, except for mass 3 ions tuned to the center of the detector resonance, well within the weakly-coupled regime. That is, $\tau_i >> \tau_c$, where τ_i (= $1/\gamma_c$) is the damping time of the ion and τ_c is the damping

time of the coil. This will be true even for coil Q as high as 50,000. Thus we can use the approximation for the frequency shift [WEI88]

$$\begin{aligned}\Delta\omega_z &= \frac{\gamma_c\gamma_i}{4}\mathrm{Re}\left(\frac{1}{\omega_0-\omega_z-i\gamma_0/2}\right)\\ &= \frac{\gamma_c\gamma_i}{4}\frac{\delta}{\delta^2+\frac{\gamma_0^2}{4}}\end{aligned} \tag{5-1}$$

where $\delta = \omega_0 - \omega_z$, γ_i is the ion's damping, γ_c is the coil damping, γ_0 is, in our approximation, nearly equal to γ_i, and ω_0 is the detector's resonant frequency. Thus the frequency shift of the ion has the shape of a dispersion curve, with the maximum shift at one half-width off the detector resonance,

$$\Delta\omega_{z,\max}(\delta) = \Delta\omega_z\left(\frac{\gamma_0}{2}\right) = \frac{\gamma_c}{4}. \tag{5-2}$$

The experiments described in this thesis were all done with $\omega_Z \approx \omega_0$, so we can approximate

$$\Delta\omega_z \approx \frac{\gamma_c}{\gamma_0}\delta. \tag{5-3}$$

Thus, if a mass 28 ion is tuned off the detector peak by, say, 250 mHz (three times more than we would expect, but much less than γ_0), the axial frequency shift will be 5 mHz for a typical Q of 32,000 and ion damping time of 3.5 seconds. This would lead to a shift in the free-space cyclotron frequency ω_c of 3.7 parts in 10^{11}, where I have used Equation 1-6 to find ω_c.

This axial frequency shift due to mistuning scales inversely with mass, since γ_Z scales inversely with mass. However, the effect of an axial frequency shift on the determination of the free-space cyclotron frequency goes as

$$\frac{\delta\omega_c}{\omega_c} \propto \left(\frac{\omega_z}{\omega_c}\right)^2\frac{\delta\omega_z}{\omega_z}, \tag{5-4}$$

so the final error due to frequency pulling scales directly as the mass, since ω_z is fixed but $\omega_c \sim 1/m$.

2. Different bad-ion cross-sections

One particularly insidious possibility is that one species of ion has systematically more bad ions than the other. Since we believe we drive all the bad ions out, that translates to one species having an occasional bad ion and the other never having any. It might be, for instance, that atomic carbon is harder to drive out than atomic nitrogen, so that CO^+ is more likely than N_2^+ to have a bad ion.

Our assurance that this is not a problem is based on the fact that we spend a lot of time adjusting the killing parameters for each ion. We are thus quite certain that, for most of the measurements, all of the bad ions are driven out. Then if one bad ion occasionally slips through the killing procedure, the distribution of cyclotron frequencies (or mass ratios) will be bimodal, or at least non-Gaussian. If one ion is actually more likely to have a bad ion, there will be a skewed distribution of outlying mass ratios. As I will show in Chapter 6, the histogram of ratios is reasonably Gaussian, so any shifts due to occasional bad ions must be too small to affect the final value. The ability to plot a histogram and compare it to a normal distribution is one clear advantage to taking the amount of data that we now can.

We do not assign an error to this effect, but subsume it into the statistical error. As long as most of the measurements do not have bad ions (which we believe to be the case), a small number of contaminated measurements will simply increase our standard deviation.

3. Differential phase error

Our method for measuring the cyclotron frequency has the possibility of introducing a fixed phase error at the $\tau = 0$ end of the phase-vs-time line. This error

doesn't change during the course of the measurement, and can be different for the two ions. Here is how the problem arises.

As described in Section III.B.2, we fix the $\tau = 0$ end of the phase line by using a direct measurement of the relative phase of the signal generators, which we calibrate to the ion phase by a number of measurements. The initial phase measurement is synchronized to the 1 Hz clock, just as the ion measurements are. However, since the two ions have different cyclotron frequencies, we must be careful to measure the initial phase and the ion phase at the same time after the 1 Hz pulse. Looking at Figure 3-5, that means we need to measure φ_{HP} at the same time (with respect to the 1 Hz pulse) as the driving pulse occurs.

In addition, we need to extrapolate the detected phase of the axial motion back to the end of the π-pulse; that is, we want the phase D mS before the data-taking starts, where D is defined in Figure 3-5. Neither this nor the timing of the φ_{HP} measurement is conceptually difficult, but there are numerous opportunities to make a mistake. The initial phase measurement must be done correctly, or an error will be introduced (much as when checking a large telescope mirror for the proper shape).

If we haven't measured the phase correction equally for both ions, we can simply keep track of one correction for each ion. However, that adds a "systematic" uncertainty due to the random errors in the phase correction measurement, as follows. Since we are measuring each ion's correction independently, there will be a natural scatter in each correction. The error resulting from this will be "locked in" for the entire measurement. Thus we need to take plenty of data when measuring the phase corrections, to ensure that the error in measurement is small compared to the final random error in the phase.

The uncertainty in an individual phase measurement is about 10°, and for the mass 28 comparison we made about 50 measurements of the initial phase for each ion. Thus we know the initial phase to about 1.5°, which is a fractional error of 2×10^{-11} in ω_c.

This error is independent of mass, and can be reduced either by taking more data for the initial phase or by carefully setting up the timing so that we can use a single φ_{HP} measurement for both ions. If we can do the latter, this error will essentially disappear.

4. Patch effect shifts

If the ions have different equilibrium positions in the trap, they will see different magnetic fields, leading to a systematic difference in cyclotron frequency. And in fact the ions do have different equilibrium positions, due to the effect of charge patches on the electrodes. This effect is negligible for mass doublets (see Section VI.G), but can be very large for non-doublets. We had measured the ratio $M(N_2^+)/M(N^+)$ (a "calibrated" non-doublet) in 1989, and found a large error in the measured value. We attributed this error to the patch effect shifts. We made another measurement of that ratio, after attempting to correct for the patch effect; this measurement is described below in Section V.D. The result was no better than before, because we can cancel the patch effect only in the $\hat{z}$ direction.

In the axial direction, however, we can measure the patch effect voltage reasonably well, and we do find that the offset is independent of ion mass. We determine the offset by adjusting the voltage on the lower endcap and looking at how the axial frequency shifts [WEI88, Section VII.A]. The frequency has a quadratic dependence on V_{LEC}, with a minimum at the electrical center of the trap. The location of the minimum is then the patch effect offset.

Figure 4-3 shows the axial frequency shift versus V_{LEC} for N^+ and N_2^+. Both species have been fit to parabolas, and the minima are in almost the same location (the error bars just touch). From this data, we can also extract some information about the electrostatics of the trap. Brown and Gabrielse [BRG86] define the odd orders of the trap potential as B_i (the even orders are labeled C_i, such as the quartic anharmonic term C_4). This convention is somewhat confusing, since we also use B_1, etc. to describe the

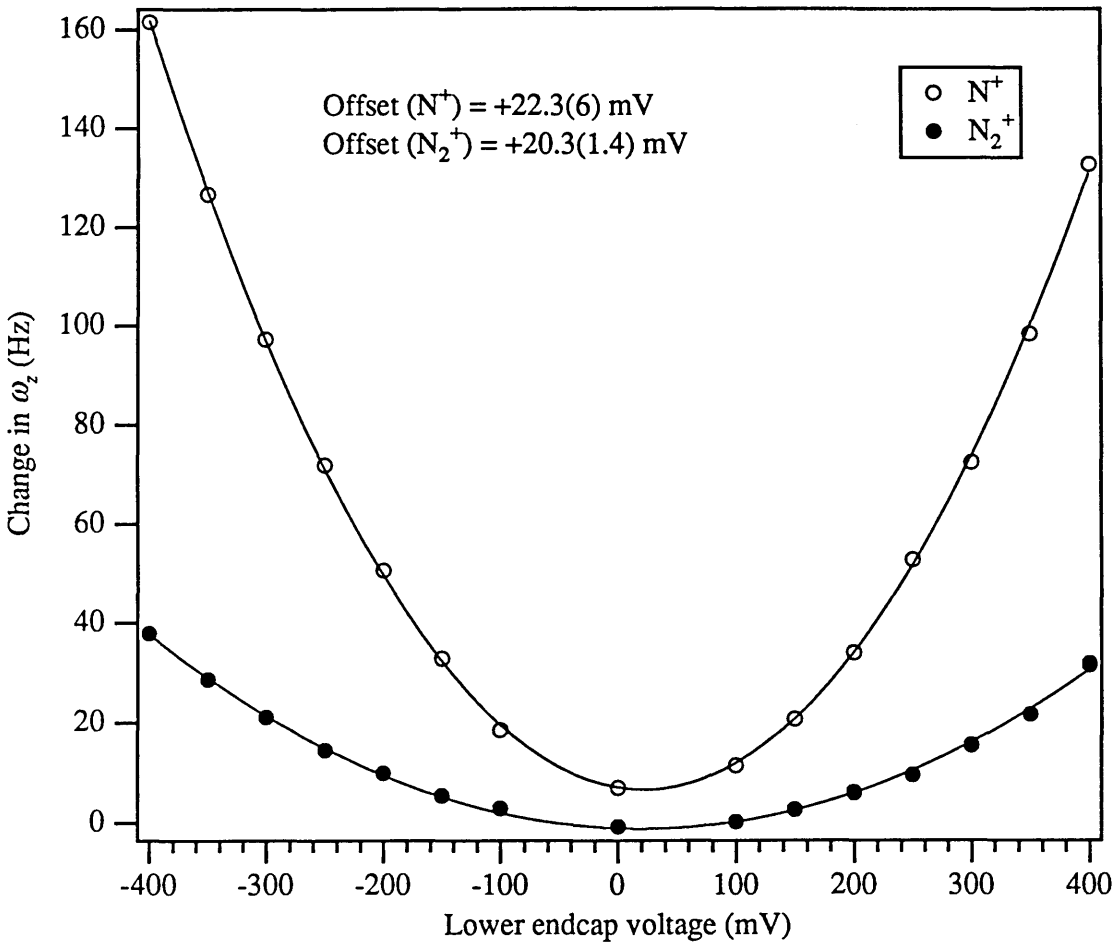

Figure 4-3. Shift in ω_z versus voltage on lower endcap. The minimum of the parabola indicates the electrical center of the trap, and thus the patch effect offset. The N^+ data have been shifted up for clarity.

magnetic field, but the meaning is generally clear from context. The curvature of the parabolas gives the product of B_1 and B_3, and we get about what we expect:

N_2^+	0.210(3)
N^+	0.206(1)
Calculated [BRG86]	0.21

Again the error bars just touch, but the fact that both ions are this close to the predicted value shows that there are no gross machining or assembly errors in the trap.

B. Amplitude-dependent errors

The errors due to the finite amplitudes of the trap modes can be determined by solving the ion's equation of motion. This has been done by Brown and Gabrielse [BRG86] and Cornell [COR90]. Both of these references cover this material thoroughly, so I will simply summarize the results. In the end, the dominant errors are due to the lowest even-order field imperfections and relativity, so I will consider only those terms.

1. The big matrix

It turns out that the dominant errors are all proportional to the squares of the amplitudes of the three trap modes. Writing the errors in matrix form, I take the result straight from [COR90],

$$\begin{pmatrix} \delta_c' \\ \delta_z \\ \delta_m \end{pmatrix} = \mathbf{D}' \begin{pmatrix} a_c^2 \\ a_z^2 \\ a_m^2 \end{pmatrix}, \tag{5-5}$$

where δ_c', δ_z and δ_m are the relative frequency shifts *of the free-space cyclotron frequency* due to the error induced in the measured cyclotron, axial, and magnetron frequencies respectively, and

$$\mathbf{D}' = \begin{pmatrix} \frac{-B_2}{2} - \frac{\omega_c^2}{2c^2} + \frac{3\omega_m C_4}{2\omega_c d^2} & \frac{+B_2}{2} - \frac{3\omega_m C_4}{\omega_c d^2} & \frac{-B_2}{2} - \frac{3\omega_m C_4}{\omega_c d^2} \\ \frac{-B_2}{2} & \frac{3\omega_m C_4}{2\omega_c d^2} & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{5-6}$$

B_2 is the coefficient (with units cm^{-2}) of the magnetic "bottle" inhomogeneity, which is of the form

$$\mathbf{B}_{\text{bottle}} = B_0 B_2 \left\{ \left(z^2 - \frac{\rho^2}{2} \right) \hat{z} - z\hat{\rho} \right\} \tag{5-7}$$

and C_4 is the fourth-order electric field gradient. Several terms, such as the special relativistic shift due to the axial and magnetron amplitudes, have been dropped, because they are entirely negligible. The first obvious point is that shifts in the magnetron frequency are also negligible. When working with mass 28, we typically measure ω_m to about 0.03 Hz, which contributes an error of 4×10^{-12} to ω_c. For lighter ions, the magnetron frequency is an even smaller perturbation to ω_c; a 1 Hz error in ω_m for mass three (for which ω_m ~250 Hz) causes an error in ω_c of only 1×10^{-13}.

2. Differential drive amplitude

Let us look at the first row of $\mathbf{D}'$, which determines the critical error in ω_c':

$$\begin{aligned}\delta_c' = \frac{3\omega_m C_4}{\omega_c d^2}\left(\tfrac{1}{2}a_c^2 - a_z^2 - \tfrac{1}{2}a_m^2\right) \\ + \frac{B_2}{2}\left(-a_c^2 + a_z^2 - a_m^2\right) \\ + \frac{\omega_c^2}{2c^2}\left(-a_c^2\right)\end{aligned} \tag{5-8}$$

There are contributions due to C_4 and B_2, and the relativistic mass shift $\omega_c^2/2c^2$. While we are measuring ω_c', the axial and magnetron motions are cooled. Therefore the second and third elements in the first row are small compared to the first element. Clearly, then, we need to consider only possible causes of a systematic difference in cyclotron amplitude. (We plan ultimately to put two ions in the trap simultaneously. In that case the magnetron amplitudes will *not* be small, and in fact can be a significant source of error.)

The first obvious source of such a difference is if we systematically drive the two ions with different amplitude pulses. It is in fact fairly difficult to ensure that the ions see the same amplitude drive, since the drives are at different frequencies. We can calibrate the drive strength quite accurately for the π-pulse (see Section III.A.2), but the drive frequency ω_c' is 160 kHz away from the coupling frequency (ω_c' - ω_z). Since the cyclotron

frequencies of the two ions differ by at most a few kHz, the π-pulse calibration does not accurately reflect the calibration at the drive frequency.

However, we were able to use the π-pulse calibration to check the transfer function of the guard ring drive cryoelectronics, at least for mass 28. We did this by calibrating the rf drive voltage entering the Dewar. When we first compared the π-pulse calibrations of CO^+ and N_2^+, we found that there was a significant difference in the strength V of the avoided crossing, even though the signal generator was set to the same nominal amplitude for both ions. We checked the actual drive voltage versus frequency (using an oscilloscope) and found that the coupling frequencies for CO^+ and N_2^+ fell on the steep edge of a resonance in our drive circuitry. (The signal generator has a 50 Ω output impedance.) This resonance, probably due to a reflection from the incorrectly terminated drive cable (see Section II.D), caused a difference in V of about 15% between the two ions.

We found that the ratio of the voltages at the two frequencies matched the ratio of the coupling amplitudes as determined by π-pulsing. That is, if we denote the measured voltage at the signal generator output as D_{CO^+} or $D_{N_2^+}$, and the coupling strength determined by the π-pulse measurement as V_{CO^+} or $V_{N_2^+}$, we found (to our measurement accuracy of 2%)

$$\frac{V_{CO^+}}{D_{CO^+}} = \frac{V_{N_2^+}}{D_{N_2^+}}. \tag{5-9}$$

This implies that a simple measurement of the driving voltage (when loaded by the actual trap circuitry) is proportional to the field at the center of the trap. It further implies that if we eliminate the reflections in the driving electronics, the transfer function will be quite flat.

For our CO^+/N_2^+ measurement, we lengthened the cyclotron drive cable, moving the peak of the cable resonance to about 4.5 MHz. Thus the driving and coupling

frequencies were at the top of the response curve, and thus flat to first order. Then the π-pulse calibration showed that the coupling strengths were equal within 1%. The resonance was also broad enough that the two driving frequencies were of equal power; at the drive frequency, the amplitude was down by less than 1% from the peak of the resonance. The effect of the remaining offset depends on the values of C_4 and B_2, and is discussed in Chapter 6. For our B_2 of about -0.1 G/cm, a 1% mismatch of drive amplitudes only causes an error of 4×10^{-13}. This value is independent of mass, although it depends directly on the ratio of the two ion masses. The ratio of $^3He^+$ to HD^+, for instance, is 10 times larger than the mass 28 ratio, so the error from a 1% mismatch is 4×10^{-12}.

3. Anharmonic frequency shifts

We now turn to the second row of $\mathbf{D}'$. This is the error introduced through errors in ω_z. Our measurement of ω_z comes from the axial frequency of the ion after the π-pulse, at which time the cyclotron radius is small. Thus the first element in this row contributes a smaller error than the B_2 term in the matrix element D'_{11}. The error is smaller by the ratio $a^2_{c,thermal}/a^2_{c,driven}$, where $a_{c,thermal}$ is the cyclotron amplitude after cooling and $a_{c,driven}$ is the amplitude during the pulse-and-phase measurement. In general D'_{11} is dominated by its relativistic term. In the future we may adjust B_2 to cancel the relativistic shift (*i.e.* to make ω'_c independent of a_c; in such a case the contribution from D'_{21} could be significant.

The axial radius, of course, is large during the measurement of ω_z. However, we need to know what the axial frequency was at the time when the cyclotron frequency was measured, *i.e.* with only thermal axial amplitude. Thus any differential axial frequency shift caused by the anharmonicity will lead to an error in the free-space cyclotron frequency.

How can the two ions have systematically different values for the axial frequency shifts? Again, if the amplitude of one ion is systematically smaller than the other, it will see a smaller shift. Since the axial amplitude is swapped over from the cyclotron amplitude, any imbalance in the initial drive pulse will map into an imbalance in a_z. In addition, it is possible for C_4 to be different for the two ions. This can happen if we use the two separate voltage settings on the VBOX for the two ion species. Then, since the trap is thus tuned individually for both ions, they can have slightly different values of C_4. In fact, this is what we did (foolishly) for the N_2^+/CO^+ ratio measurement described in Chapter 6. A much better approach would be to use only one setting of the VBOX, and adjust the computer-supplied offset voltage to move the ions into resonance with the detector.

For our mass 28 measurement, the major source of differential anharmonic frequency shift was the possibility of different values for C_4. If we assume the two ions had values of C_4 which differed by 5×10^{-6} (corresponding to half a mV on the guard rings, which looks noticeably mistuned), then the frequency estimation routine produces an estimate shifted by less than one mHz, which is 7×10^{-12} in ω_c'. In the future, using only one setting of the VBOX will leave only the amplitude scatter, which causes a shift an order of magnitude smaller.

4. Anharmonic phase shifts

One source of error which does not appear in the matrix $\mathbf{D}'$ is the effect of anharmonicity on the estimated phase. This is a somewhat subtle effect, because a fixed phase shift doesn't show up in the slope of the phase-versus-time plot, and so doesn't change the frequency. However, as mentioned above, we fix the $\tau = 0$ end of the phase line with respect to our signal generator phase at the beginning of the measurement. If C_4 then changes during the course of the measurement (due to thermal drifts in the V_{trap} circuitry, for instance), the shift in the estimated phase will drift as well.

How much phase shift can we expect from this effect? To find out, we generated some simulated anharmonic data and measured the change in estimated phase versus C_4. We found that a change of one mV on the guard rings caused a phase shift of about 12°, which corresponds to 1.7 parts in 10^{10} for the mass 28 measurement. One mV is about 10 times more change than we would expect from the VBOX, so this mechanism can only cause one or two degrees of phase shift.

Our final error after averaging a night's worth of data amounts to about 6°, so one or two degrees of systematic shift is tolerable. Also as in the case of anharmonic frequency shifts, the values of C_4 for the two ions can drift independently only if we use the two voltages on the VBOX, so as long as we don't do that we are essentially safe from this error.

5. Differential noise level

Surprisingly, a difference in the thermal noise of the ions can lead to a difference in the mean cyclotron frequency. Because the leading order shifts in the cyclotron frequency (due to B_2, C_4, and relativity) all grow quadratically with cyclotron amplitude, the thermal scatter not only causes scatter in the measured cyclotron frequency, it also causes a shift in the mean cyclotron frequency. Figure 4-4 shows graphically the effect a quadratic mapping has on any scatter in a measurement. The x axis represents the cyclotron radius and the y axis is the resulting frequency shift, due to the combination of relativity, B_2 and C_4. Clearly the mean of the scattered data $\bar{y}$ is not equal to the shift at the mean radius $y(\bar{x})$.

We can collect the coefficients for the three quadratic effects into one and write the cyclotron frequency shift as

$$\delta_c = K a_c^2, \tag{5-10}$$

where K is the combined coefficient. Then

$$\overline{\delta}_c = K\left(a_c^2 + \Delta^2\right), \tag{5-11}$$

and the resulting shift, in units of the shift due to the mean radius, is

$$\frac{\overline{\delta}_c - \delta_c(\overline{a}_c)}{\delta_c(\overline{a}_c)} = \left(\frac{\Delta}{\overline{a}_c}\right)^2. \tag{5-12}$$

Now, if for some reason one ion species is not as well cooled as the other, the difference in thermal scatter will show up as a systematic frequency difference.

As an example, if the scatter due to thermal noise is 20% of the cyclotron amplitude (roughly what we expect for mass 3), we expect an extra shift in the cyclotron frequency equal to 4% of the mean shift. For mass 3, we expect the mean shift (due almost entirely to relativity) to be about 2×10^{-9}, so the additional shift would be nearly 9 parts in 10^{11}. If for some reason one species of ion had 30% more thermal noise than the other, that would result in a systematic relative shift of 3×10^{-11}. (For our mass 28 measurement, the mean shift is an order of magnitude smaller and the scatter is less than 10% of the driven amplitude, so this effect is only about 3 parts in 10^{13}.)

It is easy to imagine several mechanisms which could result in one ion being hotter than the other. If we use the two voltages on the VBOX, for instance, the different op-amps may have different noise characteristics. It is also conceivable, though rather unlikely, that some source of narrowband noise might leak into the experiment and excite one species but not the other. Another, more likely, possibility is that we err slightly in the amplitude of one of the π-pulses. This could leave significant energy in the cyclotron mode, which would add randomly to the next driving pulse.

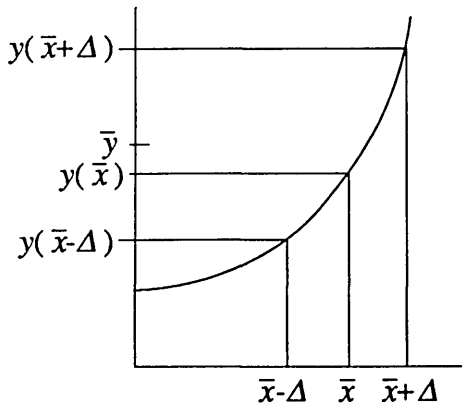

Figure 4-4. How a uniform scatter mapped through a quadratic shift causes a shift in the mean.

As mentioned earlier, we do not need to use both VBOX voltages, so that source of differential noise is not an issue. We also are quite careful in measuring the π-pulse, and always vary the pulse length, rather than the amplitude, so that nonlinearities in the driving electronics are not a problem. Moreover, our parameter estimation routine provides a fairly accurate measure of the amplitude, and we can compare the amplitudes of the two ions for each individual ratio that goes into a mass comparison. Comparing the fit amplitudes includes amplitude-dependent shifts from all sources, of course, not just this particular one. We made this comparison for our mass 28 measurement, and found a slight dependence of the mass ratio on amplitude difference. The results are presented in Section VI.E.

C. How to reduce systematic errors

The first thing to do in reducing errors is to make the magnetic field as homogeneous as possible (unless we want to use B_2 to cancel relativity, as described in [COR90]). We can directly measure only the B_1 (linear gradient) and B_2 terms of the field, but those are the most important for our work. Higher-order terms are small near the center of the trap, and radially asymmetric terms are averaged out as the ion rotates around the center of the trap. (Actually, that is not really true. We believe that the error in our measurement of $M(N_2^+)/M(N^+)$ was due to a big linear asymmetric inhomogeneity in combination with an asymmetric patch effect. See Section V.D below for details.)

1. Measuring and shimming B_1 and B_2

Our magnet has a set of 13 superconducting shim coils [FLA87] to minimize inhomogeneities in the field. Most of them have been left as initially adjusted by Oxford to provide a minimum width NMR signal using their calibration probe (which we do not have). The first- and second-order axial gradients, B_1 and B_2, however, can be measured with the ion, so we can set those terms very close to zero *in the trap*. This is important,

because the trap materials are all slightly magnetic (dia- or para-), so at the desired level of parts in 10^8, the field is perturbed by the trap.

We have measured B_1 and B_2, and shimmed them both to better than we had with trap 2, although B_2 at least is measurably nonzero still. The shimming procedure is detailed in [COR90], but I will note a couple of important details here. I will use Z_1 and Z_2 to refer to both the shim coils and the current through them; the meaning should be clear from context. Also, since we are still far from the level where we would need to use the fine shims, I refer only to the coarse shim coils.

Because there is some coupling between each shim coil and each gradient, we write the effect of the shims in matrix form,

$$\begin{pmatrix} B_1 \\ B_2 \end{pmatrix} = \begin{pmatrix} B_{1_0} \\ B_{2_0} \end{pmatrix} + \begin{pmatrix} \frac{\partial B_1}{\partial Z_1} & \frac{\partial B_1}{\partial Z_2} \\ \frac{\partial B_2}{\partial Z_1} & \frac{\partial B_2}{\partial Z_2} \end{pmatrix} \begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix}, \tag{5-13}$$

where the first term is the "natural" inhomogeneity, due to imperfections in the main solenoid and the effect of magnetic materials in and near the trap. We would like to know all these matrix elements, but as we shall see below, we do not yet have enough data for that.

It is important to note that we measure the two coefficients in very different ways. The bottle field B_2 causes a shift in the *axial* frequency, ω_z, which is proportional to the square of the cyclotron radius a_c. Thus we can measure B_2 by pulsing the ion to successively larger cyclotron radii, and measuring the resulting shift in ω_z. We determine B_1, on the other hand, by measuring the cyclotron frequency as we shift the position of the ion with respect to the magnet. Thus the accuracy with which we can determine B_1 depends on the stability of the magnetic field.

Also, when measuring B_1, we can either shift the ion with respect to the trap (by applying an antisymmetric potential to the endcaps) or move the entire trap assembly

vertically in the cryobore (by turning the hold-down screws on the top of the cryogenic insert). If the "mechanical" value differs significantly from the "electrical" value, the difference must be due to a large gradient caused by the trap or some other part of the nearby apparatus. Fortunately, we saw no such difference when we compared the two methods.

Table 5-1 summarizes the sequence of shimming attempts for trap 3, and the resulting values of B_1 and B_2. When we first recharged the magnet after the most recent quench, we reset the shim coils to their pre-quench values. (Actually, as mentioned in Section II.H, we charged the main coil backwards, so we also reversed the current in all the shim coils.) For the first shimming attempt, we tried to measure the inhomogeneities with the demountable leads still connected, so that we could adjust the shim currents iteratively. This was quite successful for the Z_2 coil, but not for Z_1.

The reason for this is that the act of shimming the magnet generates huge (parts in 10^6) shifts in the field strength B_0, due to the additional heat load on the magnet's helium bath. As mentioned above, we measure B_1 by looking at how the cyclotron frequency shifts with position in the trap. The large temporal drift in B_0 swamps the changes we see as we move the axial position of the ion (a relatively slow process). During the first shimming, we were able to see that we had moved the Z_1 current in the correct direction, but we had to wait until the next day to get a quantitative value for B_1.

After measuring B_1, we found that we were twice as far from zero, though on the other side, as when we had started. This rather large error could have been due to either a change in the strength of the Z_1 coil or a large B_1 term from the Z_2 coil. Since we could not tell from the first round of shimming which it was (it turns out that the cross-coupling from Z_2 is very small), we just set the Z_1 current to a random round number (0.00) the next time. This allowed us to determine the strength of the Z_1 coil, so that we could get close to the correct value after the third shimming.

Z_1 (A)	Z_2 (A)	B_1 (G/cm)	B_2 (mG/V^2)	B_2 (G/cm^2)	Comments
-3.64	+7.00	-0.75(4)	-7.9(3)	-0.70	Original settings
	+3.80		-12.3(5)		
	+10.50		-1.51(4)	-0.13	
-6.64		+1.51(2)	-1.02(2)		End of first round of shimming
0.00		-1.7(1)	-1.05(5)	-0.09	After second round
		-2.5(2)	-0.98(9)		Trap lowered by 7.75 mm
-5.00		-0.23(2)	-1.08(4)		
-5.48			-1	-0.09	

Table 5-1. History of shim currents and resulting inhomogeneities in Trap 3. Units have been chosen so that the exponent is 0. B_2 is given with respect to a drive voltage (applied for 40 mS), and in terms of G/cm^2, assuming that a 1 mV pulse results in a cyclotron radius ρ_c = 1.06 μm.

Before we shimmed a third time, we were able to determine the cross-coupling between the Z_2 coil and the B_1 gradient, by the method I will describe below. This kind of cross-coupling can be due to imperfect winding of the Z_2 coil, but it will also be caused by the trap not being at the field center. The value we found for the change in gradient versus change in Z_2 current was

$$\frac{\partial B_1}{\partial Z_2} = 0.24(3)\frac{G}{cm \cdot A}. \tag{5-14}$$

This implied that the trap was about 0.8 cm above the field center. (Knowing this value exactly requires knowing the strength of the Z_2 coil, which we assumed was as stated in the magnet specifications.) At that point we realized that we had not accounted for shrinkage of the insert when we positioned the trap on the insert (at room temperature). We machined a new spacer for the room-temperature end of the insert, as described in Section II.I, which lowered the trap[8] by 0.775 cm. As can be seen from Table 5-1, this had no appreciable effect on B_2, although it caused a noticeable change in B_1. It is not possible

[8]The last two digits of that value are arbitrary, but accurate.

to calculate the expected change in B_1, because we don't know the actual value of B_2 outside the trap (where the perturbations due to the trap itself are completely different).

We were not able to reduce the B_2 term below about 0.1 G/cm^2, because we ran into the maximum allowable shim coil limit. In fact, it is now set to +10.50 amps, 5% above the maximum rating. To reduce B_2 further, we will have to add some magnetic material to the trap, as we did with trap 2 [COR90].

We would like to find all four matrix elements between the first and second-order coils and their associated gradients. Unfortunately, however, the data we have do not allow us to find the change in either gradient with respect to changes in Z_2. To find the diagonal element, $\frac{\partial B_2}{\partial Z_2}$, requires an independent calibration of amplitude, which we don't have. To measure the effect on the linear gradient, $\frac{\partial B_1}{\partial Z_2}$, we would need to deliberately change Z_2, wait a day, measure B_1, and reset Z_2. Because of the (small but nonzero) danger of quenching the magnet, or even breaking it, every time we shim, we are loath to shim twice just to find out that coefficient.

We can use Table 5-1 to determine the coupling between Z_1 and both gradients, with the results

$$\begin{aligned} \frac{\partial B_1}{\partial Z_1} &= -0.47(3)\frac{G}{cm \cdot A} \\ \frac{\partial B_2}{\partial Z_1} &= 0.00(1)\frac{G}{cm^2 \cdot A}. \end{aligned} \tag{5-15}$$

The first of these is twice the value given in the magnet specifications. We do not understand this discrepancy, especially as we found a value that matched the specs exactly when we did the same measurement with Trap 2. We assume it is due to the shifting of the magnet which occurred during the last quench (see Section II.H). Such a large change in a magnet parameter is distressing, however, and the magnet should be properly shimmed at the next opportunity.

If we assume that the manufacturer's specification for $\frac{\partial B_2}{\partial Z_2}$ is still correct (not a particularly safe assumption, considering the above), we can extract a calibration for the ion's orbit size. I will do this in Section V.C.2.a below.

2. Absolute amplitude calibration

Correcting for the amplitude-dependent shifts due to the C_4 and B_2 terms does not require an absolute calibration of the cyclotron amplitude. The final perturbation, as well as the measurements of C_4 and B_2, can simply be measured in terms of the drive amplitude. Then the calibration from drive level to ρ_c cancels out. However, the perturbation due to relativity does not depend on any parameters we measure, so we must know ρ_c absolutely to determine the effect of relativity.

It is important to note that an absolute measurement of any of the three amplitudes would suffice, since the π-pulse converts amplitude from radial to axial modes in a known way [CWB90a]. For example,

$$\frac{\rho_z(\text{after } \pi\text{-pulse})}{\rho_c(\text{before } \pi\text{-pulse})} = \sqrt{\frac{\omega_c}{\omega_z}}. \tag{5-16}$$

Thus, if we know the ratio of detected signal to axial orbit size, we can immediately calibrate the sizes of the radial orbits.

a. Three old methods

We had tried three approaches to calibrate the amplitude of the ion. The first method [WEI88] was a direct measurement of the transfer function from the endcap current to the computer output. This involved combining the gains of several stages: the tuned circuit, the SQUID, the mixers and filters after the SQUID, and the computer software. While straightforward in concept, this method is difficult to implement, mainly because of the uncertainty in the parameters of the tuned circuit and the problems of generating a known 10^{-15} A current at 160 kHz.

The second approach was to excite the ion with a known drive level, and measure the resulting signal. The difficulty with this method was that the transfer functions of the cryofilters had to be measured at room temperature (where we could put a scope probe on them), and they changed significantly when cooled to liquid helium temperature. Also, the driving voltages always appear across a pair of trap electrodes, so a true differential scope input (which we did not have) is required to accurately measure the drive level even at room temperature.

The third approach, described in [COR90], was to calibrate the cyclotron amplitude using the strength of the B_2 shim coil in our magnet. The basic idea was as follows: we measured B_2, in terms of cyclotron drive strength (see table 5-1). We then changed B_2 by adjusting the current in the Z_2 shim coil. This gives us the strength of the Z_2 coil in terms of cyclotron drive strength,

$$\frac{\partial B_2}{\partial Z_2} = +1.70(6)\frac{\mathrm{mG}}{\mathrm{V}^2 \cdot \mathrm{A}}, \tag{5-17}$$

for a 40 mS pulse. Since the sensitivity of the Z_2 coil is given in the magnet specifications,

$$\frac{\partial B_2}{\partial Z_2} = +0.15\frac{\mathrm{G}}{\mathrm{cm}^2 \cdot \mathrm{A}}, \tag{5-18}$$

we can compare the two numbers to find the ratio of mV-sec to cm,

$$\text{Drive strength} = 0.106(4)\frac{\mathrm{cm}}{\mathrm{V}}. \tag{5-19}$$

This third method has the advantage that it does not rely on any electronic parameters remaining constant from room temperature to 4.2 K. However, it does assume that the magnet shim strengths given by Oxford Instruments are correct. When we did this measurement in 1990, we believed the specified shim strengths, because our calibration check of the Z_1 coil matched the calibration given in the magnet specs. We do have an absolute calibration for the displacement of the ion due to an antisymmetric voltage

across the endcaps (it's a simple electrostatics problem—see [WEI88]), so we could verify that the published strength of the Z_1 coil was correct.

Unfortunately, as mentioned above in Section V.C.1, when we redid this experiment in 1992 we found a factor of two difference between the specified strength and the measured strength of the Z_1 coil. We assume this is due to the displacement of the magnet that occurred when the magnet quenched most recently. In any case, we no longer trust the shim coil calibrations supplied by Oxford Instruments.

b. Using relativity

We have contrived a fourth method to calibrate the ion's amplitude, using the relativistic shift. The idea is to compare the shifts of the axial and cyclotron frequencies as the cyclotron amplitude is varied. Then, as we can see from (5-6), the shift in the axial frequency has only a B_2 term, while the shift in the cyclotron frequency has terms due to both B_2 and relativity. Thus we can eliminate the effect due to B_2 (assuming C_4 is small), leaving a frequency shift proportional to the square of the drive amplitude, where the constant of proportionality is known. To be precise, we need the expressions for the axial and cyclotron frequency shifts. Considering only the shifts due to the cyclotron radius, we have [COR90][9]:

$$\frac{\delta\omega_c'}{\omega_c'} = \left(\frac{-B_2}{2} - \frac{\omega_c^2}{2c^2} + \frac{3\omega_m C_4}{2\omega_c d^2}\right) a_c^2 \tag{5-20}$$

and

$$\frac{\delta\omega_z}{\omega_z} = \left(\frac{\omega_c B_2}{4\omega_m} - \frac{\omega_c^2}{4c^2} - \frac{3C_4}{2d^2}\right) a_c^2. \tag{5-21}$$

Okay, there is a term due to relativity in the axial shift, but the B_2 term is enhanced by the factor ω_c/ω_m, which is greater than 10^5 for mass 3 ions (we need to use light ions, as I

[9] Actually, the relativistic term in (5-21)` is wrong in Eq. 3.9 of [COR90] (the 4 in the denominator is written as 2. See instead [WEI88]or [BRG86]

will show below). The C_4 term is also negligible for light ions: B_2 is about 10^{-6}, so the first term is greater than 10^{-2}. Thus with a C_4 of 10^{-4} (easy to achieve), the third term is negligible. The axial frequency shift, then, is just

$$\frac{\delta\omega_z}{\omega_z} = \frac{\omega_c B_2}{4\omega_m}(Kv_c)^2, \tag{5-22}$$

where v_c is the strength, in mV-sec, of the cyclotron drive pulse, and K is the conversion between drive strength and resulting cyclotron radius, $a_c = Kv_c$. This gives us the value of B_2 in units of G/(mV-sec)^2.

Looking back at (5-20), we see that the C_4 term is again small compared to the B_2 term, so that

$$\frac{\delta\omega_c'}{\omega_c'} = \left(\frac{-B_2}{2} - \frac{\omega_c^2}{2c^2}\right)(Kv_c)^2. \tag{5-23}$$

Plugging in B_2 from (5-22) gives

$$\frac{\delta\omega_c'}{\omega_c'} = \frac{-\delta\omega_z}{\omega_z}\left(\frac{2\omega_m}{\omega_c}\right) - \frac{\omega_c^2}{2c^2}(Kv_c)^2, \tag{5-24}$$

which we can solve for K to get the desired calibration constant:

$$K^2 = \frac{1}{v_c^2}\left(\frac{\delta\omega_c'}{\omega_c'} + \frac{\delta\omega_z}{\omega_z}\frac{2\omega_m}{\omega_c}\right)\left(\frac{-2c^2}{\omega_c^2}\right). \tag{5-25}$$

One requirement for this method is that the relativistic shift must be much larger than the bottle shift:

$$\frac{\omega_c^2}{c^2} > B_2. \tag{5-26}$$

If this condition is not satisfied, the right-hand side of (5-24) is dominated by the first term, which does not depend on K. Then the value of K from Equation (5-25) consists mostly of noise from the measurement of $\delta\omega_z/\omega_z$. After the latest round of shimming, B_2 was about 6×10^{-6}, so we need $\omega_c > 7 \times 10^7$, or $f_c > 10$ MHz. Our plan is to make this

measurement with ions of mass 3, which have f_c = 43 MHz, and so clearly meet condition (5-26). This should give us an absolute amplitude calculation with an accuracy better than ±10%.

D. N^+ vs N_2^+ measurement

One very good check of systematics is to measure a non-doublet; all errors which scale with mass will be greatly enhanced. What's more, it is easy to find a "calibrated" ratio, such as atomic versus molecular nitrogen. That ratio is nearly 2.000..., with the corrections being due to the electron mass, the ionization energy, and the binding energy. These are small corrections, so the ion mass ratio is known to better than a part in 10^{12}.

1. Expected value

Let us first calculate what value we expect. If we write the atomic and molecular ionization energies, respectively, as E_{ia} and E_{im} (in units of amu·c^2), we can write the masses as

$$M(\mathrm{N}) + E_{ia} = M(\mathrm{N}^+) + M_e \tag{5-27}$$

$$M(\mathrm{N}_2) + E_{im} = M(\mathrm{N}_2^+) + M_e \tag{5-28}$$

so the ratio we measure will be

$$R_{am} \equiv \frac{M\left(N_2^+\right)}{M\left(N^+\right)} = \frac{M(N_2) + E_{im} - M_e}{M(N) + E_{ia} - M_e}, \tag{5-29}$$

where M_e is the mass of the electron in amu.

To determine the mass of N_2, we use $M(\mathrm{N}_2) = 2M(\mathrm{N}) - E_{bind}$. Then, casting (5-29) in a form where the errors are clear, we have

$$R_{am} = \frac{2M(N) - E_{bind} + E_{im} + M_e}{M(N) + E_{ia} + M_e} = 2\left(\frac{1 + \frac{E_{im} - E_{bind} - M_e}{2M(N)}}{1 + \frac{E_{im} - M_e}{M(N)}}\right) \tag{5-30}$$

From the atomic mass table [WAA85a], we find the mass is

$$M(^{14}\mathrm{N}) = 14.00307400(3)\ \text{amu}. \tag{5-31}$$

The binding energy is [HUH79]

$$\begin{aligned} E_{bind}(\mathrm{N_2}) &= 225.94(14)\ \text{kcal/mole} \\ &= 9.7962(6)\ \text{eV} \\ &= 1.05206(6) \times 10^{-8}\ \text{amu}\cdot c^2. \end{aligned} \tag{5-32}$$

Thus the mass of neutral molecular nitrogen (not needed here, but included for completeness) is

$$M(\mathrm{N_2}) = 28.00614798(6). \tag{5-33}$$

The atomic and molecular ionization energies are

$$E_{ia} = 14.534\ \text{eV [MOO70]} \tag{5-34}$$

$$E_{im} = 15.576\ \text{eV [LOF60]} \tag{5-35}$$

So the ratio we hope to find is

$$R_{am,exact} = 2.000\,039\,175\,426(1) \tag{5-36}$$

2. Data

The largest source of error when measuring a non-doublet is, we believe, the different equilibrium positions of the ions. This difference comes about because of antisymmetric charge patches. The equilibrium axial distance of the ion from the center of the trap [WEI88] is proportional to V_A/V_T, where V_T is the trap voltage and V_A is the anti-

symmetric potential. V_A can arise from voltage offsets in the VBOX or wires to the trap, from charge patches on the electrodes (by far the largest contribution), or we can apply it deliberately.

Since the charge distribution due to patch effects does not change when we change the trap voltage (see Section V.A.4), the ratio V_A/V_T, and hence the ion's position, is inversely proportional to V_T. A typical value for V_A is 20 mV, which displaces a mass 28 ion (V_T ~10 V) by 3.5 μm. Atomic nitrogen requires half as much trap voltage, hence the ion is shifted twice as far. Thus the equilibrium positions of the two ions, for a 20 mV antisymmetric patch effect, are 3.5 μm apart. A linear magnetic gradient B_1 of 0.1 G/cm would then cause a shift of 4 parts in 10^{10} between the ions.

As noted above, we can apply an antisymmetric voltage deliberately, to cancel out the axial patch effect (the procedure for measuring the strength of the patch is detailed in [WEI88], and see Section V.A.4). However, we have no control over any lateral patches which may exist in the trap. Initially, this was not much of a problem, since the field should have been radially symmetric to a high degree. (The magnet was shimmed by the manufacturer to better than 2×10^{-8}, and the trap is axially symmetric, so it can't introduce substantial radial asymmetries in the field.) However, since the magnet moved enough to crush its superinsulation during the last quench, there is every reason to suspect quite large lateral gradients in the field.

When we made the new measurement of R_{am}, we did adjust the lower endcap voltage to compensate for the measured axial patch effect. The data were analyzed in the same way as the doublet comparison described in Chapter 6, except that the cyclotron frequencies for N^+ were divided by $R_{am,exact}$ before comparing them to the N_2^+ values. Both ions should then have been at exactly the same frequency.

As can be seen from the data (Figure 5-1), the adjusted frequencies of the two ions were not at the same frequency. We have included all the measurements in this graph, although some of them are less than ideal. Most of the N_2^+ values (marked with squares)

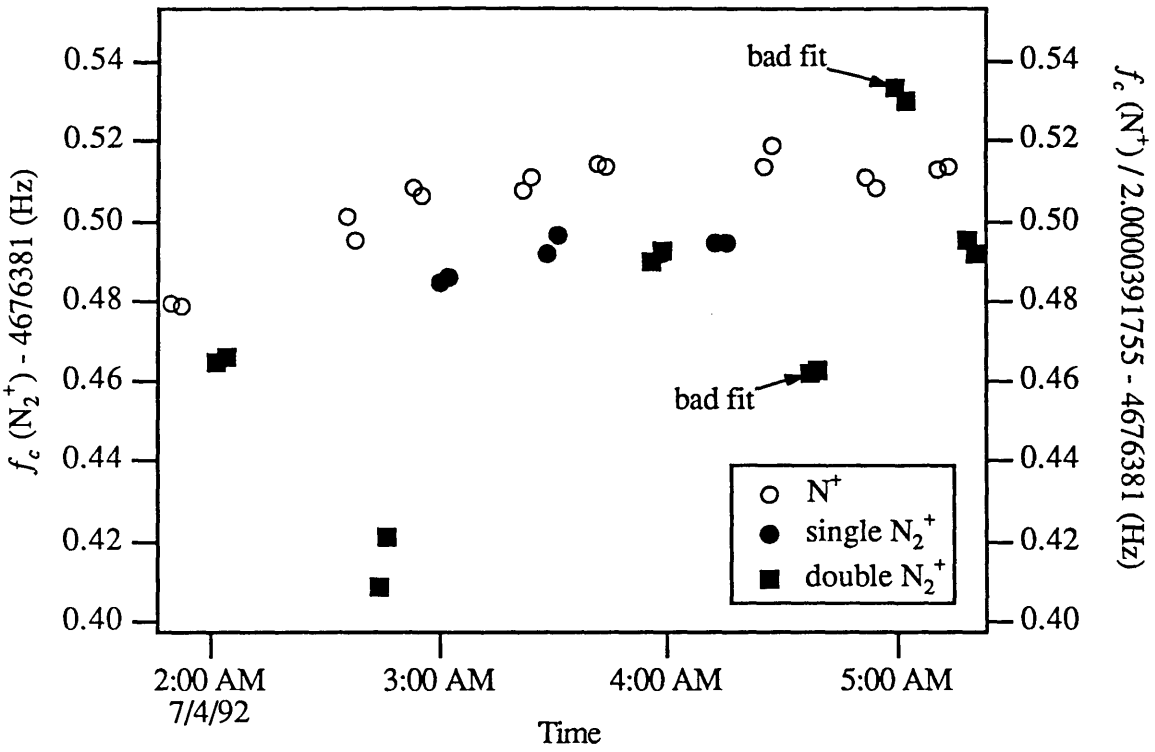

Figure 5-1. Free-space cyclotron frequencies of atomic and molecular nitrogen. The points marked "bad fit" did not give believable residuals after unwrapping the phase.

were taken with two ions in the trap, rather than one (due to an error in calibrating our automatic single-ion making routine). Also, the 7th and 8th pairs of N_2^+ numbers had excessively high residuals when the phase was unwrapped. We have included them on the graph (labeled "bad fit") for completeness.

Ignoring the "bad fit" measurements, and also dropping the one remaining outlier, we get a value of

$$R_{am,measured} = 2.000\,039\,184\,8(10). \qquad (5\text{-}37)$$

This is $4.7(5) \times 10^{-9}$ different from the correct value. We attribute this to lateral patch effects in combination with lateral field gradients. For reasonable patch potentials (~50 mV), we would need a lateral linear gradient of about 0.2 G/cm. This would result in a

VI. NEW MEASUREMENT OF $M(CO^+)/M(N_2^+)$

We have made an improved measurement of the ratio of carbon monoxide to molecular nitrogen, with the result

$$\frac{M(CO^+)}{M(N_2^+)} = 0.999\,598\,887\,74(9). \tag{6-1}$$

This is a factor of four improvement in accuracy over our previous measurement, and well inside the old error bars. This particular mass doublet (M = 28) was chosen because N_2^+ is our canonical ion, being cheap and requiring a convenient voltage ($V_T \approx$ 10V) to be resonant with our detector. Our primary source of noise, as before, was due to magnetic field variations.

This measurement was the first checkout of the rebuilt experiment, and demonstrates the result of the improvements we have made; we got four times the accuracy in about one-third the time. The data was taken on one night, and involved 33 ion switches (from CO^+ to N_2^+ or the reverse). In comparison, with our old setup we were able to make only 2 switches (CO^+ - N_2^+ - CO^+) in a night. Along with the standard improvement of $\sqrt{N}$ due to more data points, the much higher density of measurements in time allows us to see (and account for) substantially higher frequency magnetic field noise. Also, with this many data points, one can trust standard statistical techniques. Indeed, we found a Gaussian distribution of frequency ratios. Thus we no longer felt the need to (rather arbitrarily) double our statistically generated error bar, as we did before [CWB89].

The basic measurement technique was unchanged, but our ion-changing and data-taking times were greatly reduced. The typical ion-swapping time was under 5 minutes, and a pair of measurements of one ion's cyclotron frequency took even less time (see Chapter 3).

homogeneity of about a part in 10^6 over a cubic centimeter, two orders of magnitude worse than the magnet specifications. However, the magnet did move a macroscopic amount (~1 cm) during the last quench, and it has not been reshimmed since.

Because of the obvious problem of patch effect shifts, we have not attempted to make a better measurement of this non-doublet. We are considering schemes which would allow us to accumulate cyclotron phase with the same trap voltage for both ions. This should eliminate the patch effect problem, allowing us to see other (hopefully smaller) systematic errors.

Before making the measurement, we tuned the trap as well as possible for each ion, using the two sets of voltages available from the VBOX. As mentioned above, it would have been better to use one voltage from the VBOX, and adjust the computer-supplied offset voltage instead, to prevent differential drift of the axial frequency and C_4 values. We also determined each ion's cyclotron coupling frequency to within 0.1 Hz using the avoided crossing method. We began taking data after the magnetic field became quiet, as judged by eyeballing the magnetometer reading. The nearby elevator was held on the ground floor throughout the night.

A. Axial scatter

The axial frequencies of the two ions are plotted in Figure 5-2. The average standard deviation for each measurement (each contains five points) is 15 mHz (just about one part in 10^7) for both N_2^+ and CO^+. This corresponds to about 1×10^{-10} (shown by the error bar) in the free-space cyclotron frequency ω_c. Notice that the frequencies drift with respect to each other as well as absolutely. This is due to our use of both voltages on the VBOX. As we will see below, the point-to-point scatter in the frequency ratios is many times larger than the part in 10^{10} due to the axial scatter. Thus we find that magnetic field fluctuations dominate our noise (as they have in the past).

B. Extracting cyclotron frequencies

The measurement on each ion consisted of 5 phase measurements with phase accumulation times T ranging from .250 seconds to 40.200 seconds. The phase was unwrapped by sequentially estimating the frequency for measurements with successively longer T, starting from the shortest. We examined the final residuals by eye, to make sure the unwrapping algorithm hadn't gotten confused. The initial phase was calibrated as described in Section III.B.2, using the extrapolated φ_0 from all the frequency

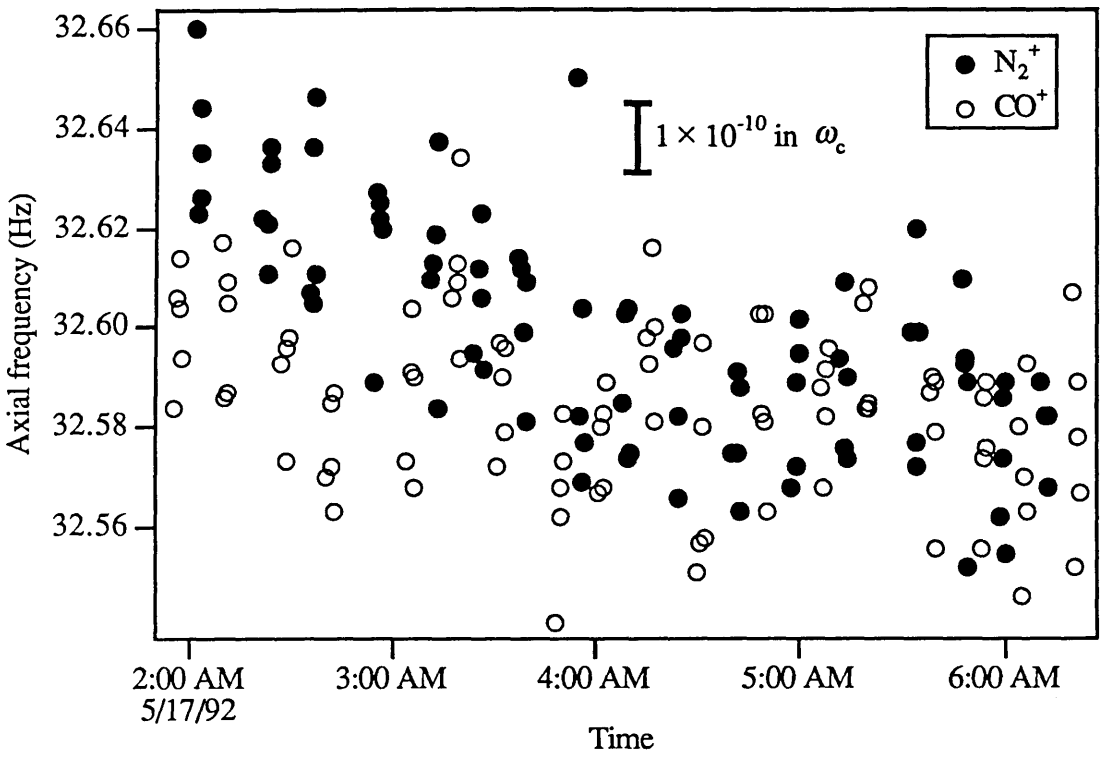

Figure 5-2. Axial frequencies of the mass 28 measurement. The error bar shows the shift that would result in a change of 1×10^{-10} in the deduced free-space cyclotron frequency.

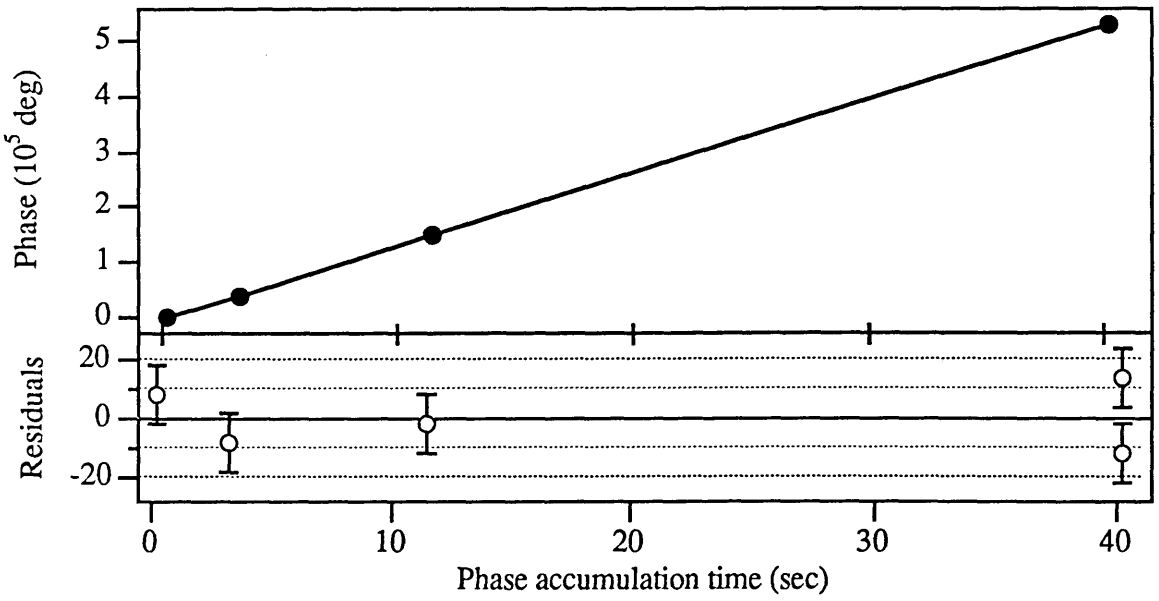

Figure 5-3. A typical cyclotron frequency measurement. The upper graph shows phase vs. accumulation time, and the lower graph shows the residuals in degrees.

measurements. Finally, ω_c' was calculated using the initial phase and the phase of the 40.200 second point. An example is shown in Figure 5-3.

C. Free-space frequencies and quadratic fit

We calculated the free-space cyclotron frequency from the three trap frequencies using (1-6). The result is shown in Figure 5-4. As one can see from the figure, the frequency drifted by about 300 mHz (7×10^{-8}) over the course of the measurement. This was a larger drift than normal, due most likely to the fact that we had shimmed the magnet two days earlier, and the field had not finished stabilizing. However, it is clear from the figure that the drift is mostly linear, and can thus be removed. The fit lines on the graph are a "quadratic plus gap" fit to all the data. That is, the fit has four free pa-

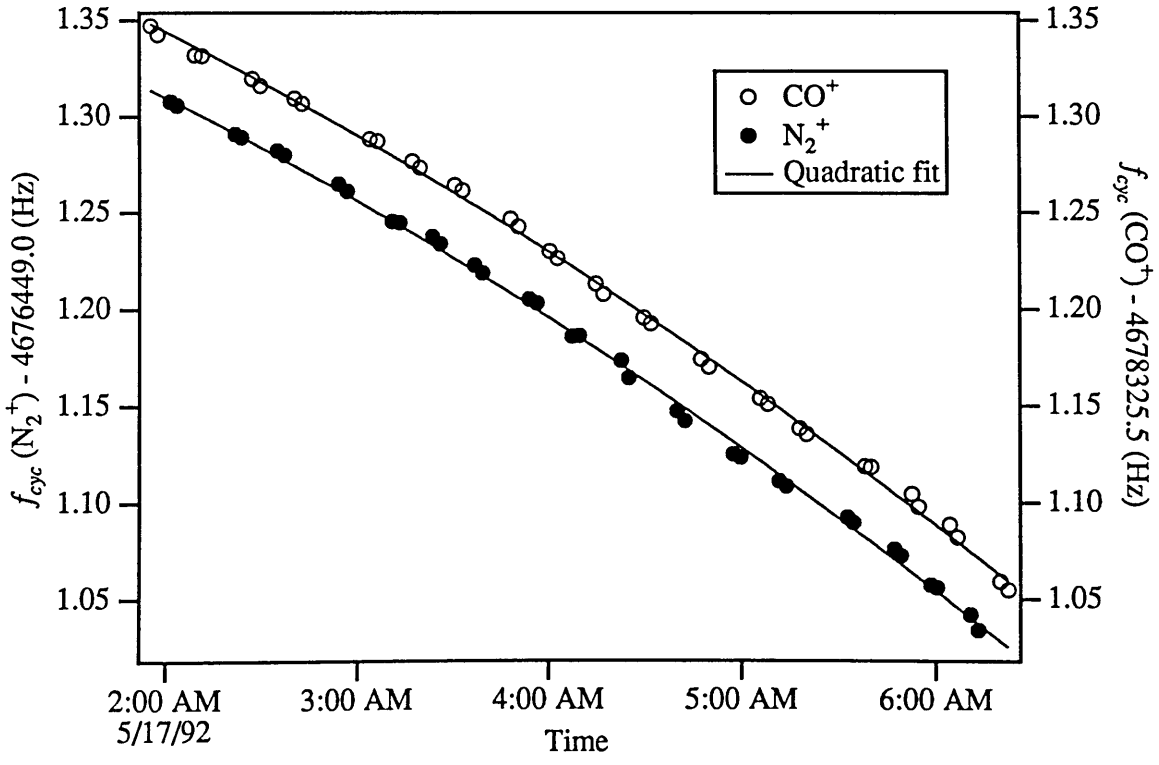

Figure 5-4. The free-space frequencies of CO^+ and N_2^+ vs time. The lines are a four-parameter fit to all the data. The parameters are constant, linear, and quadratic terms, and a fixed gap between the ions.

rameters: offset, linear and quadratic terms, and the difference between N_2^+ and CO^+.

We wanted to measure the frequency ratios of adjacent measurements, but clearly we needed to remove the long-term drift first. Otherwise, the time delay between measurements would have changed the ratios, causing the switches from CO^+ to N_2^+ yield a larger ratio than the true value, and the N_2^+ to CO^+ switches to given ratios which were too small. After averaging, that should not have affected the mean, but the calculated standard deviation would have been artificially increased. Most of the ratios would be either too large or too small, leading to a bimodal distribution.

In fact, careful examination of Figure 5-4 shows that there is higher-order structure which shows up in both ions. This may be due to fluctuations in the external magnetic field, changes in the magnet itself, or (what we consider most likely) thermal fluctuations in the support structure for the trap, causing the trap to move in the field gradient. Whatever the cause, the effect is the same as the long-term drift; the individual ratio measurements become bimodal, and the calculated error becomes unnecessarily large. Fortunately, our data points were dense enough in time that we could simply fit the drift with a higher-order polynomial.

The simple way to remove the drift is to subtract the fit of the temporal field variations from the data, leaving residuals which we then use for the measurement. However, we are looking for ratios, not differences. Thus we need to consider the error introduced by shifting the frequencies of the two ions by the same amount. If the maximum displacement due to the drift is small, we can simply subtract the fit and add back the offset term before taking ratios. If the two frequencies are f_1 and f_2, the relative error induced by a displacement H is

$$\begin{aligned}\frac{\Delta R}{R} &= \frac{\dfrac{f_1}{f_2} - \dfrac{f_1 - H}{f_2 - H}}{\dfrac{f_1}{f_2}} \\ &= 1 - \left(1 - \frac{H}{f_1}\right)\left(1 - \frac{H}{f_2}\right)^{-1}\end{aligned} \tag{6-2}$$

which, since $H << f_1 f_2$, we can approximate as

$$\begin{aligned} \frac{\Delta R}{R} &\approx 1 - \left(1 - \frac{H}{f_1}\right)\left(1 + \frac{H}{f_2}\right) \\ &\approx H\frac{(f_1 - f_2)}{f_1 f_2}. \end{aligned} \tag{6-3}$$

The total drift in this measurement was about 0.3 Hz, which results in a shift of about 2.5×10^{-11} for the last points if we remove the drift by subtraction. The shift in the final ratio is about half that, and hence negligible. However, if the precision of the experiment is to be improved by a factor of three, this error will be unacceptable.

There are two ways around this problem. One is to fit the frequencies using a ratio rather than a difference, and use that ratio as the final result. The other solution is to just remove the derivative of the fit function for each pair of measurements. Either way should eliminate this problem, and one or the other should be used for more precise measurements.

D. High-order fit and ratios

We fit the data to polynomials of increasing order, to determine what would best fit the observed drift. The fit improved quantitatively (measured by χ^2), as well as looking better in a graph, until we reached about 11th order. Above 14th order, the fitting program[10] began running into "zero pivots", indicating that the data did not support such a high-order model. Figure 5-5 shows fits to three different orders of polynomial. The constant gap from the quadratic fit has been subtracted from the CO^+ data, to bring the two ions nearly together. This way, we fit all the data with a single function. Clearly, 7th order "cuts the corners," and 13th order overshoots, whereas 11th order is close to optimal.

---

[10]Supplied by Igor, based on the polynomial fitting routine in *Numerical Recipes in C* [PFT88].

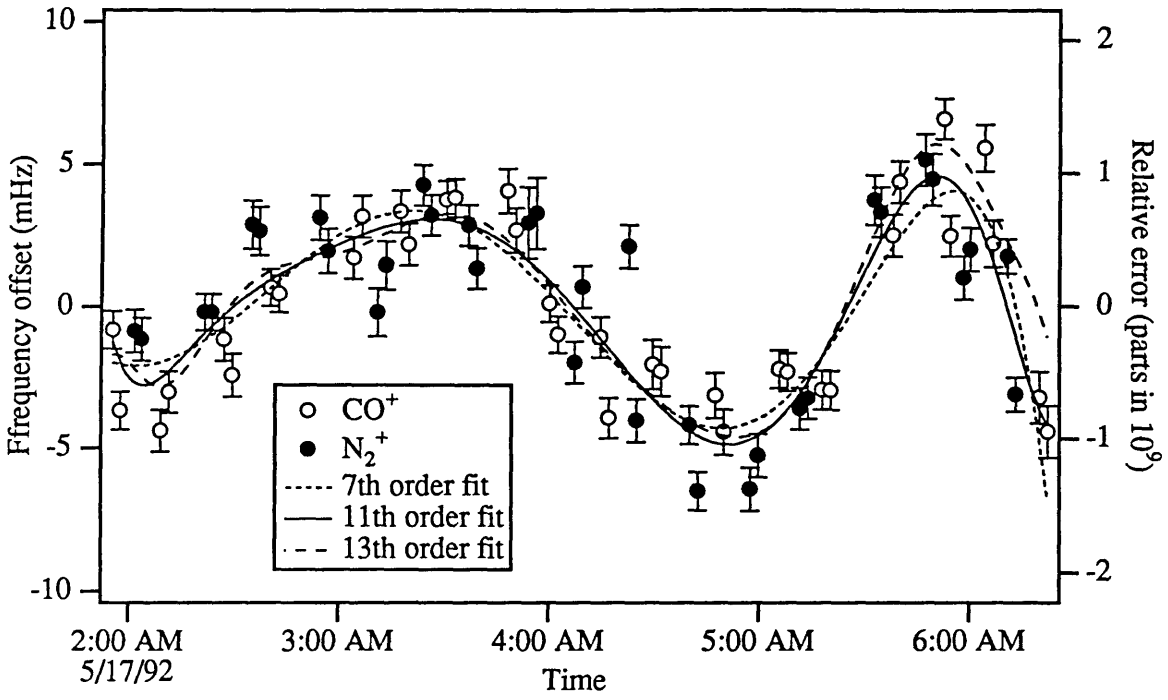

Figure 5-5. High-order fits to the data. Linear and quadratic terms have already been removed.

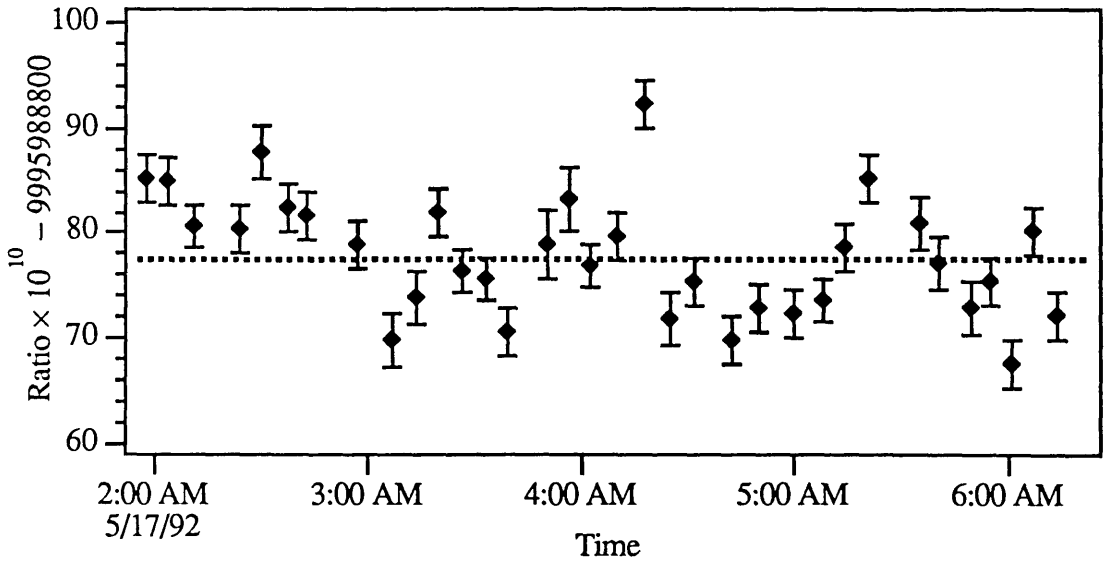

Figure 5-6. The individual ratios, plotted against the time they were taken. The dashed line is the mean value.

After removing the 11th order fit, we took the ratios of adjacent points. The results are shown in Figure 5-6. There does appear to be some structure vs. time, which may be due to some time-dependent variation of the frequency ratio which we have not accounted for. Such a time dependence of the frequency ratio would be troubling, because if we didn't know the source of it, we could not safely assume that it had zero mean. However, if it were due to the anharmonic phase shifts mentioned above, the mean would be zero. The reason for this is that the initial phase is determined by a sum of the initial phases from all the measurements (plus a batch of measurements at the beginning).

The most likely answer is that the structure in the ratios is simply due to random fluctuations. This seems probable because the "unstructured" parts of the run have a similar amount of scatter to the parts with structure. This implies that the uncorrelated fluctuations are reduced when the slow drift is at an extremum, which seems highly unlikely. In addition, the measurements appear to be distributed normally, as shown below.

Figure 5-7 shows a histogram of the ratios. For comparison, the plot also includes a Gaussian distribution with a standard deviation equal to that of the data, and a total area equal to the area of the histogram. The data appears to be distributed normally, (with the exception of one outlying point), although there are not really enough points to determine the distribution conclusively. There are 34 data points with a standard deviation of 5.7, so the error is $5.7/\sqrt{34} = 1.0$ parts in 10^{10}.

E. Dependence on detected amplitude

We looked for a correlation between frequency ratio and detected amplitude. As can be seen in Figure 5-8., there is a slight correlation. The amplitude has been squared in this plot, since we expect any frequency shifts to be quadratic with amplitude. Such a shift thus shows up as the linear correlation seen on this graph. The slope is $0.014(3) \times 10^{-10}\ mV^{-2}$.

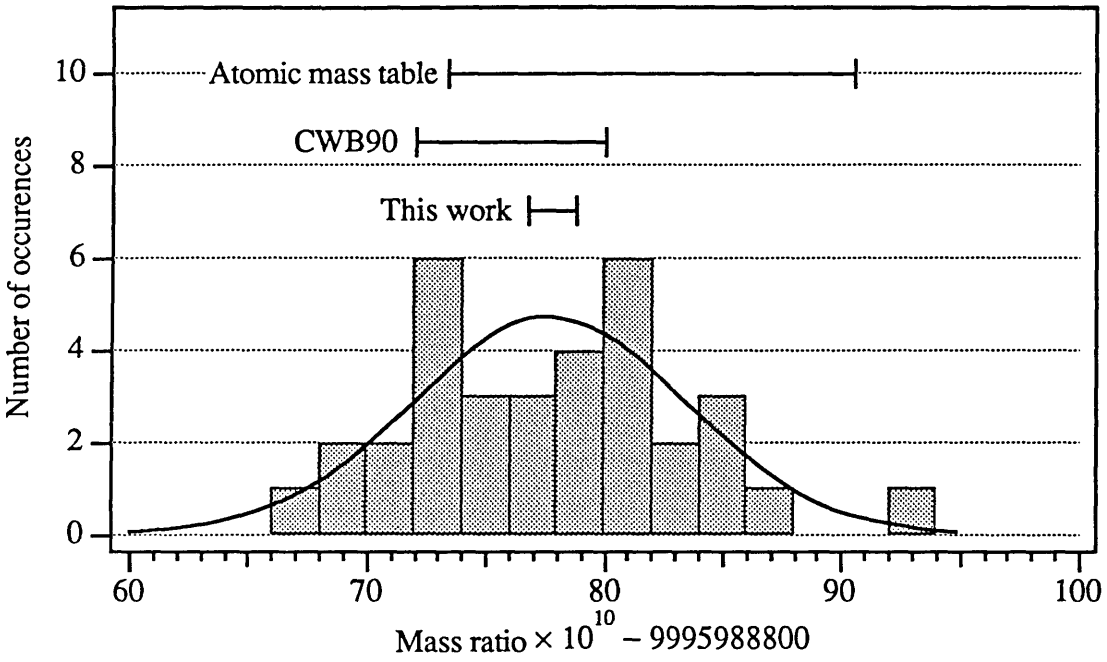

Figure 5-7. Histogram of the mass ratios. There are 34 points, and the standard deviation is 5.7. The error bars at the top show the values from the 1984 atomic mass table, our 1990 measurement [CWB89], and this measurement.

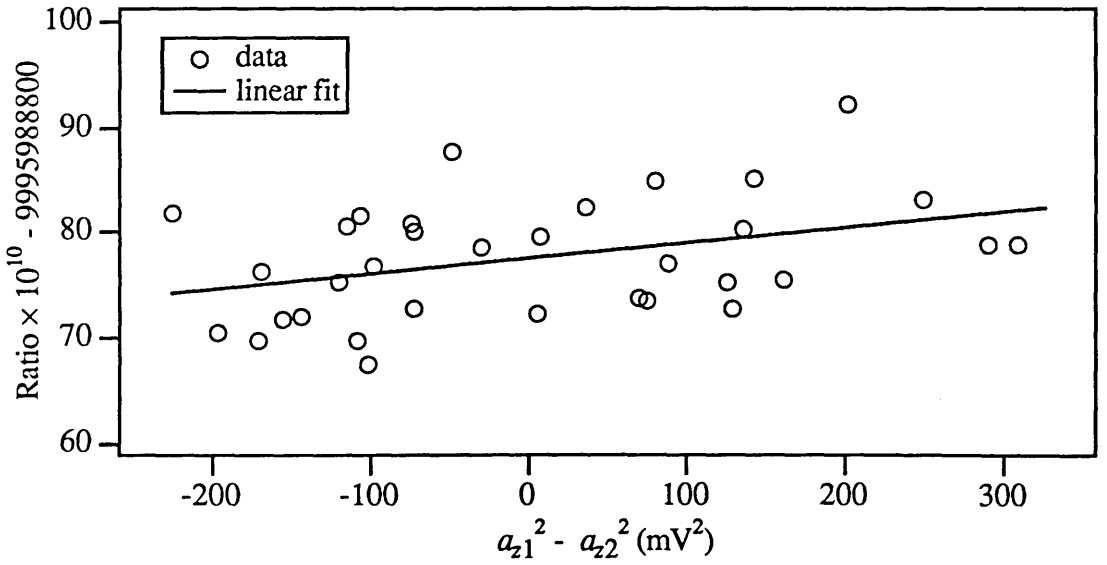

Figure 5-8. Mass ratio versus detected amplitude. The two frequency measurements which are divided to form one ratio each have a detected amplitude. These two amplitudes are squared and subtracted to get the x coordinate of each point.

We saw the same correlation when we analyzed the data using the axial frequency from the individual phase measurements, rather than the from the average of the 5 phase measurements made on each ion. Thus we know that this correlation is not due to shifts in the axial frequency caused by anharmonicity. So it must be due to some effect of the amplitude on either the cyclotron frequency or the detected phase.

The two largest amplitude-dependent shifts of the cyclotron frequency are relativity and B_2. However, they are both too small to account for the measured slope, and they almost cancel. We can use the B_2 value from Table 5-1 to find the expected slope due to the bottle shift. The table shows B_2 = -1.0 mG/V^2, but that is for the *drive* voltage, not the detected voltage. To convert, we take the average detected ("estimated") voltage for all the data points, V_{eavg} = 0.022(2) V, and set that equal to the drive voltage, V_d = 0.170(1) V. (Then, in terms of drive voltage, the slope is 2.5(5) × 10^{-8} V^{-2}.) The slope due to B_2 should be

$$\begin{aligned}\delta_c)_{B_2} &= \frac{-B_2\left[\frac{\mathrm{mG}}{\mathrm{V}^2}\right]}{2B_0[\mathrm{mG}]}\cdot(V_d[\mathrm{V}])^2 \\ &= \frac{1.05(5)}{2\cdot 85.25\cdot 10^6}(V_d[\mathrm{V}])^2 \\ &= 6.2(3)\cdot 10^{-9}V_d^2.\end{aligned} \tag{6-4}$$

To find the relativistic shift, we need to know the absolute amplitude of the cyclotron motion. We can use the approximate value of 0.106 cm/V given in Equation (5-19) to estimate that 170 mV corresponds to 1.8 × 10^{-2} cm. Then the shift due to relativity is

$$\begin{aligned}\delta_c)_{rel} &= \frac{-\omega_c^2}{2c^2}\left(V_d[\mathrm{V}]\cdot 0.106\left[\frac{\mathrm{cm}}{\mathrm{V}}\right]\right)^2 \\ &= -5.4\cdot 10^{-9}V_d^2.\end{aligned} \tag{6-5}$$

The relativistic shift is good only to the accuracy of our amplitude calibration, which may be off by a factor of 2 (though that would mean that we drive the ion axially to either

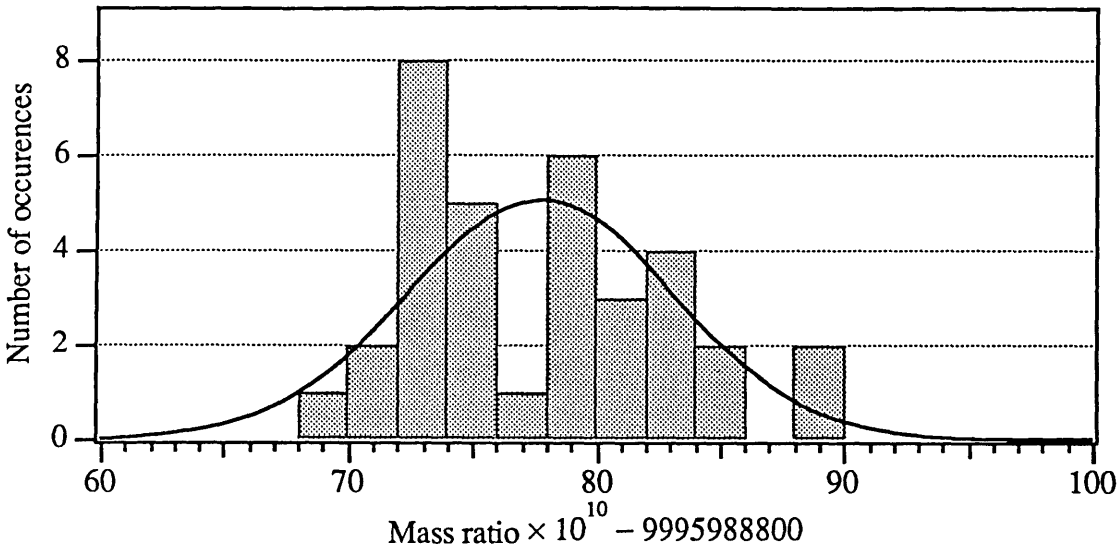

Figure 5-9. Histogram of the ratios after removing slope due to amplitude differences. The scales are the same as on Figure 5-7, and again the gaussian curve has the same area as the histogram.

10% of the trap size or 40% of the trap size, neither of which makes sense). So, since relativity has the opposite sign that the B_2 term has, these two terms can not account for the dependence on amplitude that we see.

Another possible mechanism is shifts in the estimated phase caused by anharmonicity. We have applied our estimation procedure to simulated data (with a known C_4), and measured the resulting phase shifts. The peak amplitude difference in Figure 5-8. is about 80% of the amplitude. This size of amplitude change results in a phase shift of 13°, or 1.9×10^{-10} fractional frequency shift, which is similar to the total offset on the graph. A more accurate comparison is difficult to make, since we can't exactly simulate the character of the noise, but this seems to be the best candidate to explain the dependence of ratio on amplitude.

We can correct for the slope in Figure 5-8. to further reduce the final error bars. If we remove the slope, the resulting points have a standard deviation of 5.3, which yields a final error of 0.91×10^{-10}. However, the histogram, shown in Figure 5-9, looks

somewhat less gaussian. This, combined with the fact that we have done no experiments to verify the source of this slope, makes us unwilling to quote the smaller error.

F. External magnetic field

We also want to look for any correlation between the ratios and the external magnetic field. The external field measurements are made with a fluxgate magnetometer mounted about 2 meters from the trap, where the fringing field of the main magnet is about 20 Gauss. This is above the limit of the magnetometer, so we have a pair of Helmholtz coils to bring the field in range. Although we have taken care to thermally shield the fluxgate sensor and Helmholtz coils, the measured field still drifts by a few milligauss over a couple hours. Thus we can get only short-term information on the external field.

The actual field measurements are done only during the ion's phase accumulation time, so they represent the average field during the phase measurement (plus noise). We plotted the ratios versus the difference of the field measurements, in a similar manner to Figure 5-8.. The result, shown in Figure 5-10, is that there is no detectable correlation between the ratio and the field change. This implies that internal field changes (due to pressure and temperature changes inside the apparatus) overwhelm the external field drifts. Hence, at least for our current level of precision, a self-shielding solenoid [GAT88] would not be helpful (except that it might allow us to run during the day).

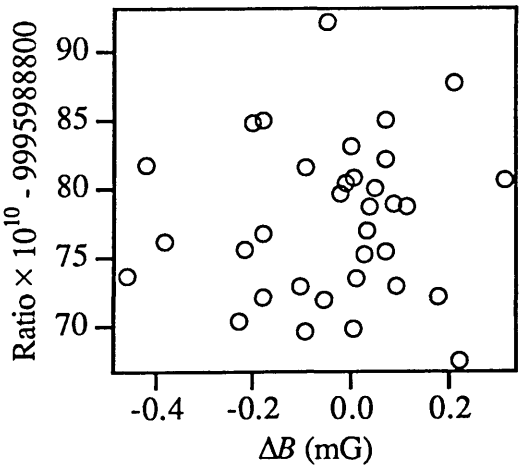

Figure 5-10. Mass ratio versus change in external magnetic field.

G. Summary of errors

We can now proceed to add up all the errors we can imagine, and hope that we have not missed any. The systematic errors are summarized in Table 6-1. The "mean shift" column shows the shift of both ions' frequencies due to each effect. The systematic difference is the level to which the shift differs between the two ions, and is what we need to asses the error in the mass measurement. The assumptions used for calculating each source of error are listed in the final column.

Some of the rows need additional explanation. Unkilled bad ions are a rare event, as mentioned in Section V.A.2, so the resulting shift in frequency can be absorbed into the statistical scatter. Therefore row 2 has been left blank; any error due to bad ions in the trap is part of the statistical error bars. The phase error rows, 3 and 7, do not lead to a net shift in the frequency, only a possible difference between the ions.

To combine the errors, we need to consider the correlations between errors. The rows which contribute are 3, 6 and 7. Rows 6 and 7 are correlated if we use the axial frequencies from the actual phase-measurement ringdowns. However, as noted earlier, we took ω_z from the average of the 5 phase measurements which constitute a cyclotron frequency measurement. This decouples the phase measurement from the frequency measurement (and, as also noted above, does not change the final result), so we can assume rows 6 and 7 are uncorrelated. Row 3 comes about by a completely independent process, so it can also be added in quadrature.

Source of error	mean shift (parts in 10^{10})	systematic difference	Assumptions
1. Tuned circuit pulling	0.1	0.05	F
2. Different bad-ion cross-sections	—	—	
3. Differential phase error	—	0.2	G
4. Patch effect shifts	0.005	0.005	E
5. Drive amplitude	0.9	0.004	A, D, H
6. Anharmonic frequency shifts	1.3	0.10	A, B, C
7. Anharmonic phase shifts	—	0.15	A, B, C
8. Noise level	0.01	0.003	A, D, I

Table 6-1. Summary of systematic errors for mass 28 measurement. The assumptions required for this table are as follows:

A. $C_4 \sim 1 \times 10^{-5}$.
B. C_4 is the same for both ions to 5×10^{-6}.
C. C_4 drifts by less than 1×10^{-6} during the measurement
D. $B_2 = -1.0(1)$ mG/V^2, and 100 μm < a_c < 250 μm, so that relativity cancels B_2 by at least 50%.
E. Lateral linear field gradient = 1 G/cm, patch effect potential = 60 mV.
F. f_z drifts by 70 mHz for N_2^+, and by 35 mHz for CO^+.
G. φ_0 measured to 1.5°.
H. Drive amplitudes equal to 1%.
I. Cooling limit of ions is 10(3)% of driven amplitude.

The final systematic error is thus 0.27×10^{-10}, which we can add in quadrature to the statistical error for a final error of 1.04×10^{-10}, which we round up to 1.1×10^{-10}. Thus the ratio is

$$R = 0.999\,598\,887\,74(11). \tag{6-6}$$

H. Mass difference

Now we need to convert the ratio to a difference. We do this by referring to (3-18), and letting ions 1 and 2 be CO^+ and N_2^+ respectively. We can use the mass of N_2 calculated in Section V.D, $M(N_2) = 28.006\,147\,98(6)$. The ionization energies are

$$E_1 = 14.013(4) \text{ eV [KRU66]}, \tag{6-7}$$

$$E_2 = 15.576 \text{ eV [LOF60]}. \tag{6-8}$$

This gives a mass difference in nano-amu (nu) of

$$M(N_2) - M(CO) = 11\ 233\ 388.8(3.1)\ \text{nu}. \qquad (6\text{-}9)$$

This can be compared to the adjusted value from the atomic mass table [WAH85],

$$M(N_2) - M(CO) = 11\ 233\ 374(26)\ \text{nu}. \qquad (6\text{-}10)$$

Thus we are well inside the error bars, as can also be seen from Figure 5-7, where the reverse conversion has been applied to convert the difference from the mass table into a ratio of ion masses.

Our result is a factor of 8 more precise than the adjusted value from the 1985 mass table. Since taking the data, we have reduced the B_1 gradient of the magnetic field by a factor of 10. This will reduce the field fluctuations seen by the ion. With three nights of data, we should now be able to make a mass comparison to 5 parts in 10^{11}.

BIBLIOGRAPHY

AUD91 Audi, G. (Personal Communication, 1991)

BKO92 Bollen, G., *et al. J. Mod. Opt.* **39**, 257-262 (1992).

BRG86 Brown, L.S. & Gabrielse, G. *Rev. Mod. Phys.* **58**, 233 (1986).

CBF92 Cornell, E.A., Boyce, K.R., Fygenson, D.L.K. & Pritchard, D.E. *Phys. Rev. A* **45**, 3049-3059 (1992).

CDB91 Camp, J.B., Darling, T.W. & Brown, R.E. *J. Appl. Phys.* **69**, 7126 - 7129 (1991).

COR90 Cornell, E.A. *Mass Spectroscopy Using Single Ion Cyclotron Resonance* (Ph. D. thesis, Massachusetts Institute of Technology, 1990).

CWB89 Cornell, E.A., Weisskoff, R.M., Boyce, K.R., Flanagan Jr., R.W., Lafyatis, G.P. & Pritchard, D.E. *Phys. Rev. Lett.* **63**, 1674-1677 (1989).

CWB90a Cornell, E.A., Weisskoff, R.M., Boyce, K.R. & Pritchard, D.E. *Phys. Rev. A* **41**, 312-315 (1990).

CWB90b Cornell, E.A., Weisskoff, R.M., Boyce, K.R., Flanagan Jr., R.W., Lafyatis, G.P. & Pritchard, D.E. *Phys. Rev. Lett.* **64**, 2099 (1990).

FHK91 Fritschi, M., Holzschuh, E., Kündig, W. & Stüssi, H. *Nuc. Phys. B* **19**, 205-214 (1991).

FLA87 Flanagan, R.W. *Trapping and Detection of Ions* (Ph. D. thesis, M.I.T., 1987).

GAB90 Gabrielse, G. *Phys. Rev. Lett.* **64**, 2098 (1990).

GAT88 Gabrielse, G. & Tan, J. *J. Appl. Phys.* **63**, 5143 (1988).

GFO90 Gabrielse, G., *et al. Phys. Rev. Lett.* **65**, 1317 (1990).

GWW90 Gerz, C., Wilsdorf, D. & Werth, G. *Z. Phys. D* **17**, 119 (1990).

HUH79 Hubert, K.P. & Herzberg, G. *Molecular Spectra and Molecular Structure Constants of Diatomic Molecules* (Van Nostrand, New York, 1979).

KAM81 Kay, S.M., and Stanley Lawrence Marple, Jr. *Proc. IEEE* **69**, 1380-1419 (1981).

KKO91 Kawakami, H., *et al. Phys. Lett. B* **256**, 105-111 (1991).

KRU66 Krupenie, P.H. *The Band Spectrum of Carbon Monoxide* National Bureau of Standards (now NIST), NSRDS-NBS-5 (1966).

KUT82 Kumaresan, R. & Tufts, D.W. *IEEE Trans. Acous., Speech, and Sig. Proc.* **ASSP-30**, 833-840 (1982).

LAL76 Landau, L.D. & Lifshitz, E.M. *Mechanics* 169pp. (Pergamon Press, Oxford, 1976).

LOF60 Lofthus, A. *The Molecular Spectrum of Nitrogen* Dept. of Physics, University of Oslo, Blindern, Norway, Spectroscopic Report-2 (1960).

MFS89a Moore, F.L., Farnham, D.L., Schwinberg, P.B. & Van Dyck, R.S. *Bull. Amer. Phys. Soc.* **34**, 99 (1989).

MFS89b Moore, F.L., Farnham, D.L., Schwinberg, P.B. & Van Dyck, R.S. *Nuc. Inst. and Meth. Phys. Rsch.* **B43**, 425 (1989).

MOO70 Moore, C.E. *Ionization Potentials and Ionization Limits Derived from the Analysis of Optical Spectroscopy* National Bureau of Standards (now NIST), NSRDS-NBS-34 (1970).

OPS75 Oppenheim, A.V. & Schafer, R.W. *Digital Signal Processing* 585pp. (Prentice-Hall, Englewood Cliffs, NJ, 1975).

RBS91 Robertson, R.G.H., Bowles, T.J., Stephenson Jr., G.J., Wark, D.L. & Wilkerson, J.F. *Phys. Rev. Lett.* **67**, 957-960 (1991).

RMC90 Rigby, K.W., Marek, D. & Chui, T.C.P. *Rev. Sci. Inst.* **61**, 834-838 (1990).

SVD81 Schwinberg, P.B., Van Dyck, R.S. & Dehmelt, H.G. *Phys. Lett. A* **81**, 119 (1981).

VFS92a Van Dyck, R.S., Farnham, D.L. & Schwinberg, P.B. *J. Mod. Opt.* **39**, 243-255 (1992).

VFS92b Van Dyck, R.S., Farnham, D.L. & Schwinberg, P.B. *High precision Penning trap mass spectroscopy of the light ions* Proc. of AMCO-9 (Bernkastel-Kues, Germany), **1**, (1992).

VMF86 Van Dyck, R.S., Moore, F.L., Farnham, D.L. & Schwinberg, P.B. *Bull. Amer. Phys. Soc.* **31**, 244 (1986).

VSD78 Van Dyck, R.S., Schwinberg, P.B. & Dehmelt, H.G. in *New Frontiers of High-Energy Physics* (ed. Korsunoglu, B., Perlmutter, A. & Scott, L.) 159 (Plenum, New York, 1978).

WAA85a Wapstra, A.H. & Audi, G. *Nuc. Phys. A* **432**, 1-54 (1985).

WAH85 Wapstra, A.H., Audi, G. & Hoekstra, R. *Nuc. Phys. A* **432**, 185-362 (1985).

WEI88 Weisskoff, R.M. *Detecting Single, Trapped Ions* (Ph. D. thesis, Massachusetts Institute of Technology, 1988).

WID74 Wineland, D.J. & Dehmelt, H.G. *Int. J. Mass Spec. Ion Proc.* **16**, 338 (1974).

WID75 Wineland, D.J. & Dehmelt, H.G. *Int. J. Mass Spec. Ion Proc.* **19**, 251 (1975).

INDEX

Page numbers for figures are in *italics*. Items in **boldface** are sections or subsections.

ACKNOWLEDGMENTS

I didn't do any of this by myself.

Dave Pritchard started the experiment, gave me the opportunity to work on it, kept the money coming in, provided expert insight, advice and editing, and made sure the helm wasn't on autopilot for too long at any one time. He also kept an encouraging face while we stumbled, Chaplin-like, into the same hole again and again.

Bob Flanagan got the bloody thing from a concept to a physical device. Robert Weisskoff singlehandedly built the original control/acquisition software, and (while finding single ions for the first time) wrote the tome which is our handbook. Anything Brown and Gabrielse didn't cover, and most of what they did, is in Robert's thesis. Eric Cornell provided many (most?) of the ideas needed to make a precision mass comparison, along with several dozen pages of utterly undocumented C code, all of which is now, thankfully, obsolete. Greg Lafyatis said (after having moved on to The Ohio State) "Why don't you warm it up now?" and we did; it didn't work again for two years. He deserves a PYI.

Vasant Natarajan trudged through a mile of knee-deep slush to meet his potential advisor in a hospital, and was still fool enough to join the group. His engineering expertise and physical insight are exceeded only by his good humor and friendliness. May his country be free of damn Brits forevermore. Happily (for me), Frank DiFilippo continues to suffer the slings and arrows of outrageous Boston. He brings a level of mathematical rigor (and a dryness of humor) to the experiment which it sorely needs. Plus he can churn out LabView code with the best of them. I wish him a rapid Ph.D. and a pennant for the Indians.

Our undergrads, Deborah Kuchnir (now Fygenson), Juan Latasa, Matt Marjanovic, Dave Kern, and now Abe Stroock, did oh so much of the grunt-work, without too much complaining.

Kris Helmerson didn't work in our sub-group, but he probably spent as much time helping people like us as he did on his own experiment. One of his jokes is still on our door.

Particularly useful help, inspiration, and levity of various kinds came from David (the renaissance man) Keith, Alex (no, I'm serious!) Martin, Chris (homebrew) Ekstrom, Wolfgang (friction $\sim v^2$) Ketterle, Rick (the lone Republican) Stoner, Michael (you call *that* a queue?) Joffe (nee Ioffe), Barbara (you used *all* my helium?) Hughey, George (Venix? No problem) Welch, and Robert (who needs Athena?) Lutwak.

Life in the lab was made bearable by the other good folks who live(d) here, including, but surely not limited to, Mike Chapman, Quentin Turchette, Troy Hammond, Brian Stewart, Warren Moskowitz, Jörg Schmiedmayer, Min Xiao, Ken Davis, Ke-Xun Sun, Scott Paine, Bruce Oldaker, Mike Kash, Marc-Oliver Mewes, Peter Martin, Pete Magill, Vanderlei Bagnato, Eric Raab, Peter Chang, Chun-ho Iu, uh...

Any such endeavor requires a large lump of support people, and there are several who were particularly supportful. Carol Costa knew (as far as it was knowable) Dave's schedule, and gracefully handled innumerable icky details of publication, travel, *et cetera*. John Peck, in RLE purchasing, went beyond the call of duty many times, and only goofed once. Gerry Power, the RLE facilities man, solved many problems (most of which were well beyond what his *job* required him to deal with), and caused none.

Quantum Devices, Inc., not only produces a SQUID sensor for half the price of BTI (even after BTI scrammed from the business) but went out of their way to bail us out no fewer than three times. They tested our (out-of-warranty) control unit and sent us a loaner rf head for three weeks, both at no charge and both with instant turnaround. They also sent a warranty replacement sensor next-day air, *before* we even sent the broken unit back.

Boston Transformer sold us a retail quantity (*i.e.* less than 250 pounds!) of square magnet wire (for the bucking coils) at wholesale prices, at our convenience, with a smile.

Wavemetrics, the makers of Igor, sent free updates without asking and always replied promptly to our questions. I don't remember Igor crashing, either.

Bill Gates, on the other hand....

Finally, my wife Kazuko not only put up with my late nights all alone with a test Dewar, but solved whatever math confusion I brought home, provided numerous helpful suggestions for this thesis, soldered the odd circuit board, and ran a mile and a half to bring the lunch I forgot for part 2 of the generals.

Thank you all.