TRIUMF Summer Work Term Report
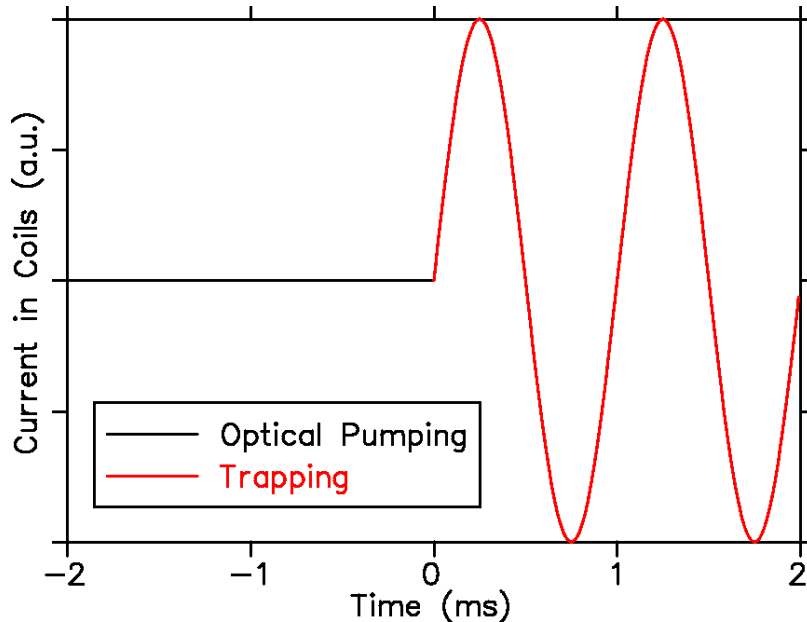
Liam Lawrence

September 1, 2016

Contents

Figure 1: Current waveform of the AC-MOT. The trapping waveform is a sinusoid.

1 Motivation and Calculations for New MOT Waveform

1.1 Background and Motivation

The TRINAT experiment captures and cools a cloud of neutral atoms using a magneto-optical trap, or MOT. The MOT at TRINAT runs current through two coils in near-Helmholtz configuration to produce the necessary quadrupole magnetic field. Experiments which require the nuclear polarization of atoms, such as recent measurements of beta decay in potassium-37, achieve spin alignment through optical pumping. During the optical pumping time, the quadrupole field must be turned off, and replaced with a dipole field. However, the atom cloud will expand during this time, and the experiment must switch back to the trapping state before the atoms stray too far. This process repeats, the apparatus alternating between optically pumping and trapping the atoms. One period of the current waveform in the main coils is shown in Figure 1. The nature of this current waveform leads to the moniker *AC-MOT* for this version of the MOT. For the remainder of this report, I will refer to this waveform as the AC-MOT waveform.

One might think a simpler waveform to use would be a square wave: off during the optical pumping time and on during the trapping time. The reason for using a sinusoid has to do with the production of eddy currents in the chamber. The main coils are surrounded by a large number of conductors, all of which become inductively coupled[1] to the main coils while the trap is running: see Figure 2 for a diagram of the TRINAT MOT.

Changes in the magnetic field induce eddy currents in nearby conductors, and these eddy currents themselves produce magnetic fields. Any off-axis magnetic field during the optical pumping time (non-dipole field) somewhat spoils the polarization of the atoms. Even if the current in the

---

[1]Recall Faraday's Law: $\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$
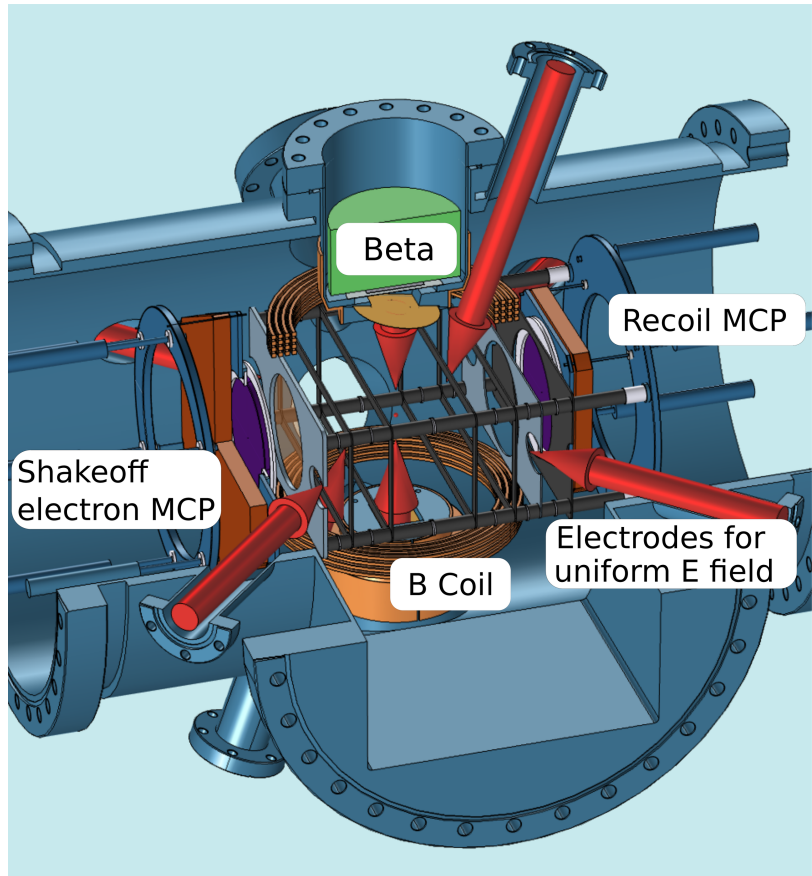
Figure 2: Diagram of the TRINAT MOT chamber.

main coils is shut off quickly (picture the falling edge of a square wave), some residual magnetic field is present since the eddy currents do not shut off instantly, but decay exponentially in time.

In a published paper, Harvey and Murray show that using several periods of a sinusoid during the trapping time mitigates the problem with the leftover field from eddy currents during the optical pumping time [4]. Melissa Anholm reproduced their results here at TRINAT: for more information, see her thesis [2].

While the AC-MOT waveform solved this problem, it introduced another: inductive heating. Since the frequency of the sinusoid is fairly high, a large portion of its power is inductively coupled to surrounding conductors, heating the metal. This is a particularly insidious problem for the delicate pellicle mirrors which are to be used in future experiments. The pellicle mirrors are mounted in a metal casing: if the mount gets too hot, the pellicle mirrors could be damaged as a result. This being a more pressing problem than the residual magnetic field during the optical pumping time, it was proposed to change the current waveform to reduce this power coupling.

Reduction of power coupling may be achieved simply by using a current waveform that results in a more slowly-varying magnetic field. Some obvious alternatives are shown in Figure 3. Before the AC-MOT waveform, TRINAT used a "square wave" produced by a device called the Russian coil driver. We model it here as a trapezoid wave because of the finite rise time. One possibility
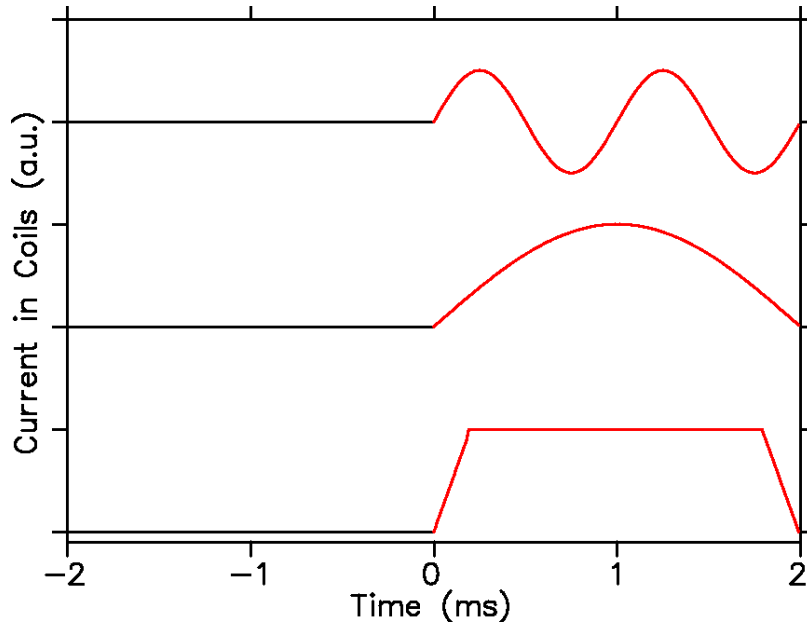
Figure 3: Alternatives to the AC-MOT waveform. Shown at the top is the original AC-MOT waveform. The middle graph is a sinusoid of a quarter the frequency. The bottom of the graph shows a trapezoid wave, such as that produced by the Russian coil driver.

was to return to this waveform. Another possibility was to use a sinusoid of a quarter the frequency to produce a waveform consisting of a single bump during the trapping time.

   The question now arises: how can we be sure these waveforms will reduce the coupled power? If they do, by how much? And which is better? To answer these questions, we turn to some mathematics.

1.2   Fourier Analysis

Taking after Harvey and Murray, we model the main coils and the surrounding conductors as coupled inductors. Provided the mutual inductance is small compared to the self-inductance of the main coils, we may assume the coupled inductors do not significantly load the main coils. In this case, we may consider a single coupled conductor individually. Furthermore, we shall assume the main coils are powered by a current generator. In the TRINAT lab, the main coils are actually voltage-driven; as a result, the voltage waveform (which the experimenter provides) is convolved with the impulse response function of the coils ($\exp(-t/\tau_{coil})/L_{coil}$ where $\tau = L_{coil}/R_{coil}$) to produce the current waveform. However, by observing the resultant current waveform (and still assuming insignificant loading), we may model the power supply as a current generator which produces the observed current waveform.

   Using the previous assumptions, we model the main coils and coupled conductor as pictured in Figure 4. The variables used are as defined in the caption.
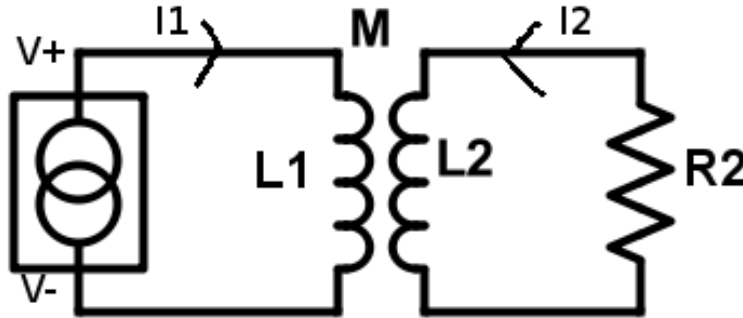
Figure 4: The equivalent circuit used to model the main coils and a coupled conductor. The main coils, of self-inductance L_1, are driven by a current generator. The coupled conductor has a self-inductance of L_2 and a resistance of R_2. The mutual inductance between the main coils and the coupled conductor is M. The current in the main coils and the coupled conductor is $i_1(t)$ and $i_2(t)$ respectively, with the directions as indicated.

The equations describing this setup are as follows:

$$M\frac{di_1}{dt} + L_2\frac{di_2}{dt} + Ri_2 = 0. \tag{1}$$

We would like to determine the average power dissipated in the coupled conductor for a given waveform in the main coils $i_1(t)$. This is most easily accomplished with a Fourier analysis approach, as will be shown. Suppose the waveforms $i_1(t)$ and $i_2(t)$ have the frequency spectra (Fourier transforms) $I_1(j\omega)$ and $I_2(j\omega)$ respectively. That is,

$$I_1(j\omega) = \int_{-\infty}^{\infty} i_1(t)\exp(-j\omega t)\,\mathrm{d}t \tag{2}$$

and similarly for $i_2(t)$. Given equation 1, we can relate I_1 to I_2 by a transfer function. Taking the Fourier transform of equation 1 reveals this fact:

$$j\omega M I_1(j\omega) + j\omega L_2 I_2(j\omega) + RI_2(j\omega) = 0 \tag{3}$$

$$j\omega\frac{M}{L_2}I_1(j\omega) + j\omega I_2(j\omega) + I_2(j\omega)/\tau = 0 \tag{4}$$

$$I_2(j\omega) = -\frac{M}{L_2}\frac{j\omega}{j\omega + 1/\tau}I_1(j\omega) \tag{5}$$

where we have defined τ as

$$\tau = \frac{L_2}{R_2}. \tag{6}$$

The mean-square value of an arbitrary power waveform $f(t)$ in the time domain is directly related to its spectrum in the frequency domain $F(j\omega)$ by Parseval's theorem, which says:

$$\lim_{T\to\infty}\frac{1}{2T}\int_{-T}^{T}|f(t)|^2\,\mathrm{d}t = \frac{1}{2\pi}\int_{-\infty}^{\infty}|F(j\omega)|^2\,\mathrm{d}\omega. \tag{7}$$

6

If the function $f(t)$ is periodic with a period $2T_0$ then $f(t)$ is expressible as a Fourier series instead of a Fourier transform. In this case, Parseval's theorem simplifies to a sum over the spectrum instead of an integral.

$$\frac{1}{2T_0} \int_0^{2T_0} |f(t)|^2 \ \mathrm{d}t = \frac{1}{2\pi} \sum_{n=-\infty}^{\infty} |c_n|^2 \tag{8}$$

where c_n is the n-th coefficient in the complex Fourier series of $f(t)$:

$$c_n = \int_{-T_0}^{T_0} f(t) \exp(-jn\omega_0 t) \ \mathrm{d}t \tag{9}$$

and $\omega_0 = 2\pi/2T_0 = \pi/T_0$ is the base angular frequency.

We shall apply Parseval's theorem for periodic functions to our current waveforms, since $i_1(t)$ will be periodic. Note we will have to multiply both sides by the resistance through which the current flows to obtain the actual power and not just the mean-square value. Let the complex Fourier series coefficients of $i_1(t)$ be c_n. Then the n-th coefficient of the complex Fourier series of $i_2(t)$, d_n, is related to those of c_n by a discretized form of equation 5:

$$d_n = -\frac{M}{L_2} \frac{jn\omega_0}{jn\omega_0 + 1/\tau} c_n \tag{10}$$

Using Parseval's theorem and the preceding coefficients, we may conclude that the average power P carried by the induced current waveform $i_2(t)$ is:

$$P = \frac{M^2 R_2}{2\pi L_2^2} \sum_{n=-\infty}^{\infty} \frac{n^2 \omega_0^2}{n^2 \omega_0^2 + \frac{1}{\tau^2}} |c_n|^2 \ . \tag{11}$$

The problem of calculating the average power dissipated in coupled conductors by an arbitrary current waveform in the main coils has reduced to the problem of evaluating the sum in equation 11 using the Fourier coefficients corresponding to the current waveform within the main coils.

1.3   Comparison of Alternatives

To compare the power dissipated in coupled conductors by the current waveform of Figure 3, we begin by determining their Fourier series. We shall label the waveforms from top to bottom by the names *high-frequency sinusoid*, *low-frequency sinusoid*, and *trapezoid wave*.

Let the period of the waveform be $2T_0$, and let the graphs of Figure 3 extend from $-T_0$ to T_0. Let the amplitude (maximum value) of the waveforms be A. Note that this means the trapezoid wave will carry approximately twice the power of the sinusoidal waveforms if they all have the same amplitude. However, since the strength of the atom trap is a function of magnetic field strength, it is waveforms of equal amplitude (not power) we wish to compare. The Fourier coefficients for one of these waveforms, $f(t)$, are given by:

$$c_n = \int_{-T_0}^{T_0} f(t) \exp(-jn\omega_0 t) \ \mathrm{d}t \tag{12}$$

7

Note that the waveforms are zero for $-T_0 < t < 0$. Also note that for the trapezoid wave we have the following notation: t_w is the full-width half-maximum of the trapezoid and t_r is the rise time (0 to maximum) of the trapezoid edge. We have made several simplifying assumptions for the trapezoid wave. We assume $t_w = T_0$, that is, the width of the trapezoid is half the period — this is a fair assumption if the rise time is small compared to the period. We also assume that the rise time and fall time are equal. See the included document **signals.pdf** for the complete expression and derivation for the trapezoid wave.

High-frequency sinusoid

$$c_n = \begin{cases} 0 & n \text{ even and } n \neq \pm 4 \\ \frac{\pm A}{4j} & n = \pm 4 \\ \frac{4A}{\pi(4^2-n^2)} & n \text{ odd} \end{cases} \tag{13}$$

Low-frequency sinusoid

$$c_n = \begin{cases} 0 & n \text{ odd and } n \neq \pm 1 \\ \frac{\pm A}{4j} & n = \pm 1 \\ \frac{A}{\pi(1-n^2)} & n \text{ even} \end{cases} \tag{14}$$

Trapezoid wave

$$c_n = \begin{cases} 0 & n \text{ even} \\ \frac{A}{2} \frac{(-1)^{(n-1)/2}}{\frac{\pi}{2}n} \frac{\sin(\frac{\pi}{2}\frac{t_r}{T_0}n)}{\frac{\pi}{2}\frac{t_r}{T_0}n} \exp(-jn\omega_0\frac{t_w+t_r}{2}) & n \text{ odd} \end{cases} \tag{15}$$

The corresponding summations for the power coupling are as follows.

High-frequency sinusoid

$$P_{\text{hfs}} = \frac{M^2A^2R_2}{2\pi L_2^2}\left(\frac{1}{8}\frac{16\omega_0^2}{16\omega_0^2 + 1/\tau^2} + \sum_{n\text{-odd}} \frac{n^2\omega_0^2}{n^2\omega_0^2 + \frac{1}{\tau^2}}\frac{16}{\pi^2(16-n^2)^2}\right) \tag{16}$$

Low-frequency sinusoid

$$P_{\text{lfs}} = \frac{M^2A^2R_2}{2\pi L_2^2}\left(\frac{1}{8}\frac{\omega_0^2}{\omega_0^2 + 1/\tau^2} + \sum_{n\text{-even}} \frac{n^2\omega_0^2}{n^2\omega_0^2 + \frac{1}{\tau^2}}\frac{1}{\pi^2(1-n^2)^2}\right) \tag{17}$$

Trapezoid wave

$$P_{\text{trap}} = \frac{M^2A^2R_2}{2\pi L_2^2}\left(\frac{1}{4}\sum_{n\text{-odd}} \frac{n^2\omega_0^2}{n^2\omega_0^2 + \frac{1}{\tau^2}}\frac{1}{\left(\frac{\pi}{2}n\right)^2}\frac{\sin^2\left(\frac{\pi}{2}\frac{t_r}{T_0}n\right)}{\left(\frac{\pi}{2}\frac{t_r}{T_0}n\right)^2}\right) \tag{18}$$

The sums in the first two equations may be evaluated using the residue theorem from complex analysis and we obtain a simpler form for the power coupling expressions. Define $a = \omega_0\tau$.

High-frequency sinusoid

$$P_{\text{hfs}} = \frac{M^2 A^2 R_2}{2\pi L_2^2} \left(\frac{4}{1 + 16a^2} - \frac{8a^3 \tanh\left(\frac{\pi}{2a}\right)}{\pi(1 + 16a^2)^2} \right) \tag{19}$$

Low-frequency sinusoid

$$P_{\text{lfs}} = \frac{M^2 A^2 R_2}{2\pi L_2^2} \left(\frac{a^2}{4(1 + a^2)} - \frac{a^3 \coth\left(\frac{\pi}{2a}\right)}{2\pi(1 + a^2)^2} \right) \tag{20}$$

The sum for the trapezoid wave, unfortunately, cannot be simplified so easily. However, it is trivial to prove that the sum converges. First, note that all terms are positive and therefore the partial sum for $-m \leq n \leq m$ is monotonically increasing for increasing m. As for the magnitude:

$$0 < \frac{n^2 \omega_0^2}{n^2 \omega_0^2 + \frac{1}{\tau^2}} \quad < \quad \frac{n^2 \omega_0^2}{n^2 \omega_0^2} = 1 \tag{21}$$

$$0 < \sin^2\left(\frac{\pi}{2} \frac{t_r}{T_0} n\right) \quad < \quad 1 \tag{22}$$

$$0 < \sum_{n\text{-odd}} \frac{n^2 \omega_0^2}{n^2 \omega_0^2 + \frac{1}{\tau^2}} \frac{1}{\left(\frac{\pi}{2}n\right)^2} \frac{\sin^2\left(\frac{\pi}{2} \frac{t_r}{T_0} n\right)}{\left(\frac{\pi}{2} \frac{t_r}{T_0} n\right)^2} \quad < \quad \sum_{n\text{-odd}} \frac{1}{\left(\frac{\pi}{2}n\right)^2} \frac{1}{\left(\frac{\pi}{2} \frac{t_r}{T_0} n\right)^2} \tag{23}$$

$$\sum_{n\text{-odd}} \frac{1}{\left(\frac{\pi}{2}n\right)^2} \frac{1}{\left(\frac{\pi}{2} \frac{t_r}{T_0} n\right)^2} \quad = \quad \frac{16 T_0^2}{\pi^4 t_r^2} \sum_{n\text{-odd}} \frac{1}{n^4} \tag{24}$$

$$= \quad \frac{16 T_0^2}{\pi^4 t_r^2} \frac{\pi^4}{96} \tag{25}$$

$$= \quad \frac{T_0^2}{6 t_r^2} \tag{26}$$

$$0 < \sum_{n\text{-odd}} \frac{n^2 \omega_0^2}{n^2 \omega_0^2 + \frac{1}{\tau^2}} \frac{1}{\left(\frac{\pi}{2}n\right)^2} \frac{\sin^2\left(\frac{\pi}{2} \frac{t_r}{T_0} n\right)}{\left(\frac{\pi}{2} \frac{t_r}{T_0} n\right)^2} \quad < \quad \frac{T_0^2}{6 t_r^2} \tag{27}$$

So long as t_r is nonzero, the sum converges. We are therefore justified in using computer code to evaluate the sum by computing the partial sum for $-m \leq n \leq m$ using an m such that adding the $m + 1$-th term produces a negligible effect on the series.

To get some concrete numbers out of the preceding math, we will need the values of relevant parameters, most of which we will have to calculate approximately. Since we will not know the accuracy of such calculations, we will take several measures to extract meaningful information. First, to reduce the number of unknown parameters, we shall compare *ratios* of coupled power. Notice that the ratios

$$\frac{P_{\text{lfs}}}{P_{\text{hfs}}} \text{ and } \frac{P_{\text{trap}}}{P_{\text{hfs}}} \tag{28}$$
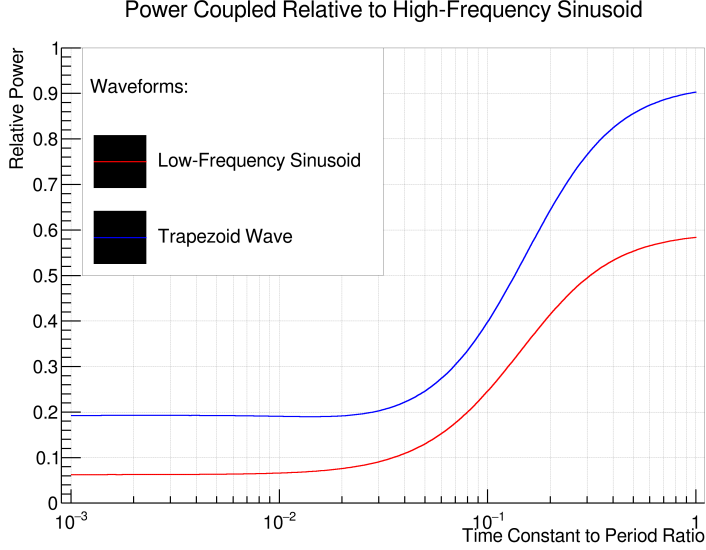
Figure 5: Plot of the ratio of the power coupled by a given waveform to the power coupled by the high-frequency sinusoid as a function of the parameter $\tau/2T_0$. For the pellicle mirror mount, $\tau/2T_0 \approx 0.0215$. At this value on the x-axis, the low-frequency sinusoid couples 12.8 times less power than the high-frequency sinusoid and 2.5 times less power than the trapezoid wave.

do not include the parameter M or A since all the power expressions contain the factor

$$\frac{M^2 A^2 R_2}{2\pi L_2^2} \tag{29}$$

which cancels out. Furthermore, we shall plot these ratios as a function of $\tau/2T_0$ since τ will not be calculable to great accuracy and we may be interested in the behaviour for different conductors (which will have different values of τ in general). It will be more informative to view the behaviour of these power ratios over a range of values of τ than at one particular value.

All that being said, let us estimate relevant parameters. Because of the previous waveform for TRINAT, I set $2T_0 = 4$ ms. Based on tests with the Russian coil driver, I estimate the rise time of the relevant trapezoid wave $t_r \approx 0.24$ ms. Using the book "Inductance Calculations" by Grover, I estimate the self-inductance of the pellicle mirror mount to be $L_2 \approx 0.044$ μH. Its resistance is $R_2 \approx 0.51$ mΩ. Therefore, the time constant of the pellicle mirror mount is $\tau \approx 86$ μs.

I have written code using C++ and Root to evaluate the power coupling sums and plot the ratios as a function of $\tau/2T_0$. The code and instructions for using it may be found in the directory **Calculations/**. Make sure to open **README.txt** first. The resulting plot is shown in Figure 5. It is easy to see both the alternative waveforms are significantly better than the high-frequency sinusoid. What is perhaps more surprising is that the low-frequency sinusoid is even better than the trapezoid wave at reducing the power coupling.

To reduce the problem of inductive heating in the TRINAT chamber, we are thus prompted to use a current waveform consisting of a single sinusoidal bump in the main coils — the middle waveform of Figure 3 which I have been calling the low-frequency sinusoid. The MOT subject

to this new waveform has been christened the *Rectified*-AC-MOT, or RAC-MOT[2]. The subject of the following two sections is how we equipped the hardware of the MOT's control system to use this new waveform, and the tests we performed to measure the effectiveness of the RAC-MOT by modifying existing software.

2 RAC-MOT Hardware

This section is a description of the hardware configuration used to implement the RAC-MOT. We begin with a brief overview of the existing hardware that is relevant to this project and move on to the updates from there.

2.1 Existing Hardware

Atom Traps There are actually two atom traps at TRINAT. There is a smaller trap which is used to extract and collect the element of interest, and a larger trap which houses all the equipment used to study the atoms. I will refer to these as the *small trap* and the *main trap* respectively. The atoms collected in the small trap are transferred to the main trap by a push-beam of laser light.

Raspberry Pi The Raspberry Pi, or the Pi as I will refer to it, is a small computer running Linux. The Pi controls and monitors all the instruments needed to run the experiment using its digital I/O pins and a control code. Claire Preston wrote this code; please see her report to learn more. Her report also explains the whole system in more detail than I will present here. Later in this report I will speak of "modifications to the software," by which I mean modifications to the code that runs on the Pi.

LRS Boxes Very close to the Pi is a set of electronic boxes I will refer to as the LRS boxes (though not all of them are made by LRS). The LRS boxes are responsible for performing Boolean logic operations and pulse timing modifications on various digital signals.

Acousto-optic modulators The acousto-optic modulators, or AOMs, modify the frequency of the laser light sent to the trap. A sinusoidal electrical signal is applied to piezoelectric material connected to an optical crystal, which produces vibrations of frequency f in the crystal. Incoming laser light is diffracted by the crystal, but since the lattice is vibrating, the light which exits at diffraction mode m has its frequency shifted by mf. The power of the electrical signal sent to the AOMs may be used to directly control the intensity of the light sent to the trap.

Function generators TRINAT has many function generators, two of which we use to produce the RAC-MOT waveform. There is a B&K Precision 4052 and a Stanford Research Systems DS345. I will refer to these function generators as the BK and the SRS respectively. Both of these function generators may be triggered by an external logic signal (the rising or falling edge of the external logic signal will produce a waveform burst at the output of the generator)

---

[2]Pronunciation is still up for debate: some say "ray-cee-mot" while other say "rack-mot."

Figure 6: The blue waveform is a measure of the (inverted) current output of the Matsusadas. The pink waveform is the voltage input to the Matsusadas. The cyan waveform is a TTL logic signal which triggers each pink sinusoid and amplitude modulates the same.

and have their output amplitude modulated by an external analog signal. Note that the SRS can be in triggered burst mode and have its output amplitude modulated simultaneously. The BK can only be in one of these states at a given time.

Matsusadas The Matsusadas are a pair of heavy-duty amplifiers which drive the main coils in the main trap. Each amp is hooked up to its own coil, but both receive the same input signal to be amplified. The Matsusadas may be placed in current output or voltage output mode. When running a DC (constant) trap, the Matsusadas are set to output approximately 56 amps.

IR Camera As part of her work term, Claire Preston set up an IR camera to take pictures of the fluorescing atoms in the trap. The camera faces directly into the centre of the trap at a vacuum window which allows light to escape.

2.2 RAC-MOT Configuration

I will explain the hardware configuration by starting simply and adding layers of complexity. A complete hardware diagram is available at the end of this section. The reader may want to read the section, take a look at the diagram, then read the section again to obtain a solid understanding.

2.2.1 How to produce the desired waveform

The goal of reconfiguring the hardware is to produce a waveform like the one shown in Figure 6, which I will refer to as the *RAC-MOT waveform*. The simplified version goes like this:

1. The Matsusadas are set to voltage output mode.

2. The SRS is set to external trigger burst mode. Upon receiving a rising edge at the trigger input, the SRS will output one period of a sine wave (usually set to a frequency of 250 Hz, or a period of 4 ms).

3. The BK is set to output a square wave which simulates a TTL signal (0 to 5 V) with an adjustable duty cycle. I will call this waveform the BK TTL. The frequency of this square wave must be less than the frequency of the SRS sine pulse (less than 250 Hz in normal operation). We generally used around 210 Hz.

4. The output of the BK is hooked up to the external trigger input of the SRS and the amplitude modulation input of the SRS.

5. When the SRS receives a rising edge of the square wave from the BK, it will output its one-period sinusoid. This sinusoid is amplitude modulated by the BK's square wave. The duty cycle of the square wave may be adjusted to set the time at which the sinusoid is shut off.

6. The output of the SRS is fed to the input of the Matsusadas

Notice in Figure 6 that the current output of the Matsusadas (blue) does not perfectly match the input waveform from the SRS. This is because the Matsusadas are in voltage control mode and they drive a significantly inductive load. The voltage output of the Matsusadas *does* look like the pink waveform of Figure 6. The current output is related to the voltage output by convolution with the impulse response function of the coils, which has the form

$$h(t) = \frac{1}{L} \exp(-Rt/L) \tag{30}$$

where L is the self-inductance of the coils and R is the resistance. This smooths out the waveform to produce the blue waveform of Figure 6.

This is the reason for having a negative overshoot on the sinusoid produced by the SRS. The duty cycle of the BK TTL signal can be adjusted to ensure the negative overshoot exists for just such a time as to make the current output of the Matsusadas zero during the optical pumping time. Otherwise, the exponential decay of the current in the Matsusuadas will cause the magnetic field to be non-zero during the optical pumping time.

The reason we do not simply use the Matsusadas in current output mode is because they appear not able to handle such an inductive load at high currents. One will see damped oscillations at the edges of a waveform which turns on or off quickly.

Another thing to note is that although in the mathematical analysis we assumed a period of 4 ms, with the trapping phase and optical pumping phase being 2 ms each, this is not how it always works out in practise. The trapping phase is more like 2.9 ms long because of the rather sluggish response of the coils, while the optical pumping phase is usually kept a bit shorter, at around to 1.9 ms long. The reason for this is made clear later in the report.

2.2.2 Timing considerations

The preceding is not the whole story, however. The main trap will slowly lose atoms because of collisions which eject atoms from the trap. We will need to regularly replenish the trap with more atoms, and for this we must alternate between catching atoms pushed over from the small trap — which requires a strong DC field — and running the waveform shown in Figure 6.

The Pi controls this process and its timing. Of particular interest to us is a logic output pulse from the A0 pin of the Pi (in the Pi code, this pin is given the name ACMOT). The pulse from the A0 pin is only a few tens of milliseconds long: it is fed into an LRS box which lengthens the pulse to be approximately 480 ms long. I will refer to this longer 480 ms pulse as the *A0 pulse*. This pulse is used to start and stop the RAC-MOT waveform; the waveform should start on the rising edge of the pulse and stop on the falling edge.

The (480 ms) A0 pulse is monitored by the Pi on an input pin. While the pulse is high, the Pi simply waits. When the pulse falls to low, the Pi knows the RAC-MOT waveform is finished, and it moves on to the next part of the control code.

The question is: how do we ensure the RAC-MOT waveform only runs while the A0 pulse is on? Look back at step 4 of the preceding list. Instead of hooking up the BK TTL directly to the external trigger input of the SRS, we run it over to an LRS box and take the logical AND of the BK TTL and the A0 pulse (in practise, we have blanked the BK TTL with (NOT A0), which amounts to the same thing). We send this logic signal, (BK TTL AND A0), to the external trigger input of the SRS. We leave the BK TTL directly amplitude modulating the SRS, however.

This way, the SRS will only be triggered while the A0 pulse is high; if the A0 pulse is low, (BK TTL AND A0) will be low and the SRS will not be triggered and will output nothing. While A0 is high, (BK TTL AND A0) is logically equivalent to BK TTL alone and the waveform will run normally. The reason we use the BK TTL output directly to amplitude modulate is because the A0 pulse falls during the final period of the RAC-MOT waveform. If we used (BK TTL AND A0) to amplitude modulate, the final period of the waveform would be messed up.

To ensure that the BK TTL signal begins precisely with the A0 signal, we set the BK in externally triggered burst mode and hook up A0 to the trigger input of the BK. The number of square wave cycles to output upon a trigger is set to be sufficient to cover the entire length of the A0 pulse. For example, if the period of the square wave is 4 ms, since the length of the A0 pulse is 480 ms, one should set the BK to output $480/4 = 120$ cycles upon a trigger.

2.2.3 Driving the AOMs

The AOMs which feed trapping laser light to the main trap must be on during the catching phase, while the small trap pushes over atoms to the main trap. During the RAC-MOT waveform, the AOMs must be on when trapping is desired (while the sinusoidal pulse is on) and off during the optical pumping time. The AOMs may be shut on and off by a set of switches which pass or block the electrical signal that drives the AOMs.

While the A0 pulse is high, the AOMs should follow the BK TTL. While the A0 pulse is low, the AOMs should be on all the time. This is accomplished by creating the following logic combination with the LRS boxes and sending it to the aforementioned switches: (BK TTL OR (NOT A0)).

Note that in practise, while the RAC-MOT is running, the AOM is not fed directly by the BK TTL; rather, the rising edges of the BK TTL trigger a logic pulse sent to the AOM switches. The
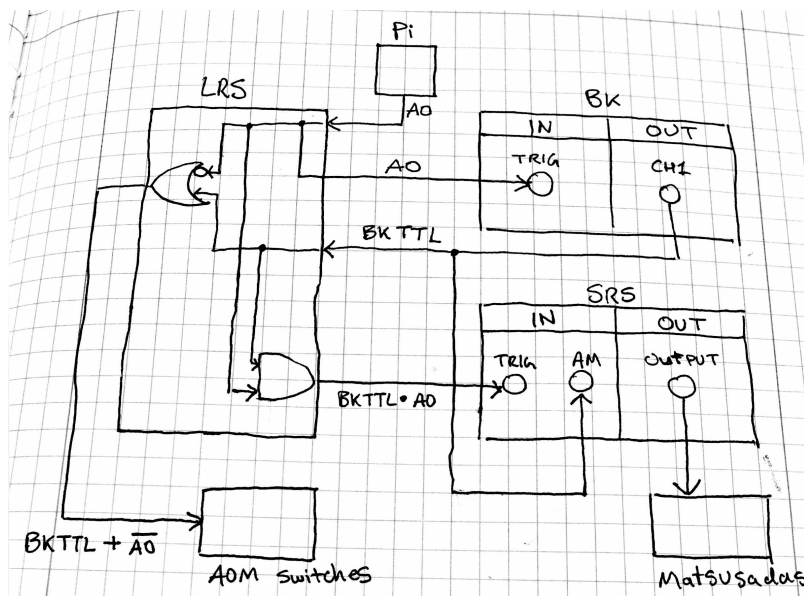
Figure 7: Diagram of the hardware setup for the RAC-MOT.

timing of this logic pulse is not set by the BK TTL, but by the LRS box which produces the logic pulse. The length must be set manually to synchronize with the BK TTL duty cycle.

2.2.4 Full diagram

See Figure 7 for a diagram which summarizes the preceding description of the hardware. I have used a shorthand notation for logic. $A \cdot B$ denotes (A AND B), $A + B$ denotes (A OR B), and $\bar{A}$ denotes (NOT A). Note that the inputs for the BK and SRS are on the back of the boxes and the outputs on the front. Also note that the signal exiting the Pi and exiting the LRS are both called A0, although technically the signal exiting the Pi is a shorter pulse. As soon as the signal enters the LRS box, assume it is lengthened to the usual 480 ms.

3 Characterizing the RAC-MOT

We tested several measures of performance to determine how well the RAC-MOT functions to trap atoms. These measurements serve as useful diagnostics: one might test potential trap improvements by repeating these measurements. In describing these tests, I will break each down into four parts: what we wanted to measure, how we measured it, how to do it yourself, and the results of the measurement. Each measurement required tweaking the experiment's software, so it is necessary to first become acquainted with the existing code.

For all these measurements, we used stable potassium-41, obtained by sublimation from a solid source. The main results of each measurement may be found in the folder **Measurements/**.

3.1 A Primer on the Software

Let us begin by examining the existing software. There are three machines to be aware of: the Pi and two laptops. One laptop, trinatlt2, controls the IR camera. The other laptop, trinatlt3, is used by the experimenter to access the Pi and trinatlt2 from one location.

When running a measurement, use trinatlt3 to access trinatlt2 and the Pi with the command `ssh`. You will want two terminals open: one for the connection to each machine.

The Pi contains the main control code under the folder **RunTrapSequence_cp**. The shell script **StartSequenceRAC.sh** will do several things when called:

1. Run the code on trinatlt2, which will capture and store a picture whenever the IR camera receives a trigger from the Pi.

2. Run the trap sequence code on the Pi.

3. Open an instance of Physica on trinatlt3 which will display the pictures stored on trinatlt2 as they are taken.

The trap sequence code on the Pi is C++ code whose name begins with **TrapSeq_cp3**. This is the code we were modifying (actually, making copies of, then modifying) to change the behaviour of the trap system to make desired measurements. One must also modify the file **params.txt** on the Pi, which contains a list of trap parameters. Of chief importance to us for these measurements were:

Max_n The **TrapSeq_cp3** code increments a variable **n** each time the main trap completes a cycle of capturing from the small trap and running the RAC-MOT waveform. This variable is used for stepping through values of some quantity, if desired, like the laser frequency. The variable **Max_n** determines the maximum value of **n** before resetting back to one.

NumPulses This variable determines how many pulses of the push-beam are used to send atoms from the small trap to the main trap. The greater this value, the more atoms will be pushed over each cycle.

ABDAC0Shift The frequency of the laser light which arrives at the main trap is determined by several AOMs which shift the frequency of the light. This variable is the voltage value output by a DAC which determines the frequency shift of one of these AOMs.

The reader should now be familiar enough with the operation of the trap to understand the following terms I will use throughout the remainder of the report. I will use the term *period* to refer to a single period of the RAC-MOT waveform (lasting 4.8 ms). The term *cycle* refers to a single iteration of catching atoms from the small trap and running the RAC-MOT waveform for its full 480 ms — in terms of the software, every time the variable **n** increments, one *cycle* has passed. A *run* consists of **Max_n** cycles, after which **n** resets to 1.

The laptop controlling the camera, trinatlt2, records the pictures taken by the camera in arrays. These arrays are saved as plain-text files with the name **cp*.out**, where the asterisk denotes the picture number. These files, as well as all the code run by trinatlt2, exists in a rather buried folder whose name I have stored in an environment variable on trinatlt2 called `fly`. Simply enter `cd $fly`

in the command line on trinatlt2 to be taken to the folder[3]. You will probably want this to be your working folder from the connected terminal on trinatlt3.

In this same folder is another file called **params.txt**. These parameters, however, are for the camera operation. The only one of relevance for our purposes is **expTimems**, which controls the exposure time of the camera.

After an experiment has been run and pictures are stored on trinatlt2, you will want to copy the pictures to another computer for storage. I have created two terminal functions for this purpose. In the `fly` folder, enter `mvdata` in the command line to move all the pictures to a sub-folder **data/** which will be created automatically (note that if this sub-folder exists already it will be destroyed). On the blue HP-Stream Notebook you are doubtlessly reading this on, use the command `yankdata` to copy the pictures from that **data/** sub-folder created on trinatlt2 to your machine (obviously, you will have to be connected to the network). `yankdata` requires one argument, a folder name. `yankdata` will create a folder of this name in the current directory and put all the pictures there. You will have to enter the password of trinatlt2 to access it.

In each section of "how to do it yourself," I will list the relevant files which must be changed to run the code. The format of these listings will be (machine):(filename). The (filename) will only be the end of the file name, not the fully directory name, since there is only one relevant directory for each machine: the ones previously mentioned. For example, **Pi:StartSequenceRAC.sh** refers to the file **/home/pi/RunTrapSequence_cp/StartSequenceRAC.sh** on the Pi.

The exception to this convention is in the naming of the Physica code used to analyze the data. The code will be referred to by **llawrence:[Physica_code].pcm**. You will find the Physica code in the folder **PhysicaCode/**. Please refer to the Physica code itself for an explanation of how to use it; I have included comments in the code.

3.2 Lifetime Measurements

3.2.1 What did we want to measure?

A MOT will lose atoms naturally because of collisions with gas particles in the imperfect vacuum, collisions with high-temperature atoms in the trap, and light assisted collisions [1]. For low densities (not too many atoms in the trap), the former two mechanisms dominate. In this regime, the number of atoms in the trap as a function of time, $n(t)$, is well-described by the following differential equation:

$$\frac{dn}{dt} = L - \lambda n. \tag{31}$$

The constant λ is the characteristic loss rate. For multiple sources of loss, each with their own rate λ_i, the total loss rate λ will be the sum of all the λ_is. The constant L is the rate at which atoms are added to the trap. If there is no loading mechanism, then $L = 0$ and the solution to the above equation is:

$$n(t) = n(0) \exp(-t/\tau) \tag{32}$$

where $\tau = \lambda^{-1}$ is the *lifetime* of the trap. The number of atoms in the trap decays exponentially in time. The longer the lifetime, the better the trap. For this reason, our first set of measurements

---

[3]The folder's full name, for reference, is:
/home/trinat/flycapture/flycapcode/src/cpAsyncTriggerEx/

consisted of lifetime measurements of the trap with the RAC-MOT waveform, to determine the trap's effectiveness. We also measured the lifetime of the trap with a DC magnetic field (always on) immediately afterwards to compare.

These lifetime measurements give us a way of optimizing the trap, as one may measure the lifetime for different trap parameter settings (e.g laser frequency) and determine the configuration of the trap which gives the greatest lifetime.

3.2.2 How did we measure it?

The basic idea is this: load the main trap with a large number of atoms (pushed over from the small trap), then take a series of pictures with the IR camera of the atoms' fluorescence while the main trap runs and the number of atoms decays exponentially in time. The total collected light from the cloud is proportional to the number of atoms in the trap, and this series of pictures therefore gives a measure of the number of atoms in the trap as a function of time.

To compare the RAC-MOT lifetime to the lifetime of a DC trap, the code will run the RAC-MOT for a given time, then switch over to a DC trap and repeat the measurement.

We use a Physica code to extract the total fluorescence as a function of time from the pictures and fit decaying exponentials. The time constant of these exponentials is the lifetime of the trap.

We repeated these measurements for various values of laser frequency, power, and optical pumping time.

3.2.3 How can you do it yourself?

Relevant files:

Pi:StartSequenceRAC.sh Shell script which begins the Pi code and the camera code.

Pi:TrapSeq_cp3_RAC_lifetimes_ABDAC1toBDC C++ code for the Pi to run.

Pi:params.txt System parameters.

trinatlt2:params.txt Camera parameters.

llawrence:lifetimes.pcm Physica code to extract RAC and DC lifetimes from the measurement.

Recipe:

1. Open StartSequenceRAC.sh and change the code called on the Pi. The instruction may be found at the bottom of the file: change it to `sudo ./TrapSeq_cp3_RAC_lifetimes_ABDAC1toBDC`.

2. Open the params.txt file on the Pi and change **NumPulses** to 3. This ensures that plenty of atoms are pushed over from the small trap to the main trap at the beginning of each lifetime measurement. The variable **Max_n** for this code determines how many cycles are run for a lifetime measurement — note that the RAC-MOT and DC-MOT are run immediately after each other, and the maximum number of cycles **Max_n** is split between them. That is, if **Max_n** is 60, then the RAC-MOT will run for 30 cycles, then the DC-MOT will run for 30 cycles, then the code repeats.

Table 1: Relationship between **ABDAC0Shift** variable and the frequency shift of the laser from the corresponding AOM, as measured in the lab. Note I have listed the frequency shift divided by two, since the laser passes through this AOM twice. The first value has been linearly interpolated from the other values.

ABDAC0Shift Value	Frequency Shift/2 (MHz)
1.24	67.5 (interpolated)
1.29	68.52
1.34	69.57
1.39	70.48
1.44	71.49

3. Open the params.txt file on trinatlt2 and change the exposure time, **expTimems** to 25.0.

4. Call `./StartSequenceRAC.sh` from the command line on the Pi to start the code.

5. The code will run indefinitely; type `q` and press Enter on trinatlt3 to stop the measurement.

6. Call `mvdata` on trinatlt2 and `yankdata [folder_name]` on the HP-Stream.

7. Enter the folder you just created with `yankdata`, where the copied pictures reside, and run Physica. Execute the code **lifetimes.pcm**.

3.2.4 What were the results?

First, it bears mentioning that for some reason the lifetime is systematically worse for the first run. After the first run, the lifetime is fairly consistent for subsequent runs.

A typical lifetime plot upon running the Physica code resembles that shown in Figure 8. We tested the dependence of the lifetimes upon: laser frequency, laser power, and optical pumping time. See Figures 9, 10, and 11 for the plots.

A few things should be said for each. The laser frequency is set by the variable **ABDAC0Shift** in the params.txt file on the Pi. I have included a table below which shows the correspondence between the value of **ABDAC0Shift** and the frequency shift divided by two. See table 1.

The laser power is controlled by a digital attenuator which sets the power in the electrical signal sent to the AOMs, which determines how much laser power is present in the first-order diffracted light (sent to the trap). The integer variable in the Pi code called **istep**, ultimately determines how much the laser power is attenuated. If **istep** is zero, the laser power is not attenuated. The laser power decreases monotonically (but not linearly) with increasing **istep**. We only took three measurements, using **istep** equal to 0, 1, and 2. This corresponds to roughly 100% power, 75% power, and 50% power respectively. At 100% power, the laser power is approximately 200 mW.

The optical pumping time is simply the time in between pulses of the RAC-MOT sinusoid. This may be adjusted by changing the frequency and duty cycle of the BK TTL. Notice that the lifetime drops dramatically when the optical pumping time is greater than about 2 ms. For this reason, the optical pumping time has been kept at around 1.9 ms, as stated previously.
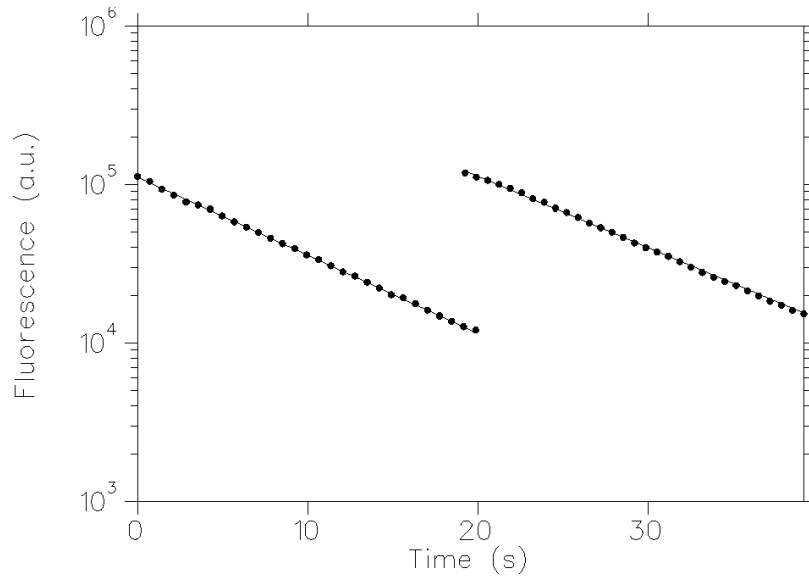
Figure 8: Typical plot of atom fluorescence versus time in a lifetime measurement with a log scale on the y-axis.
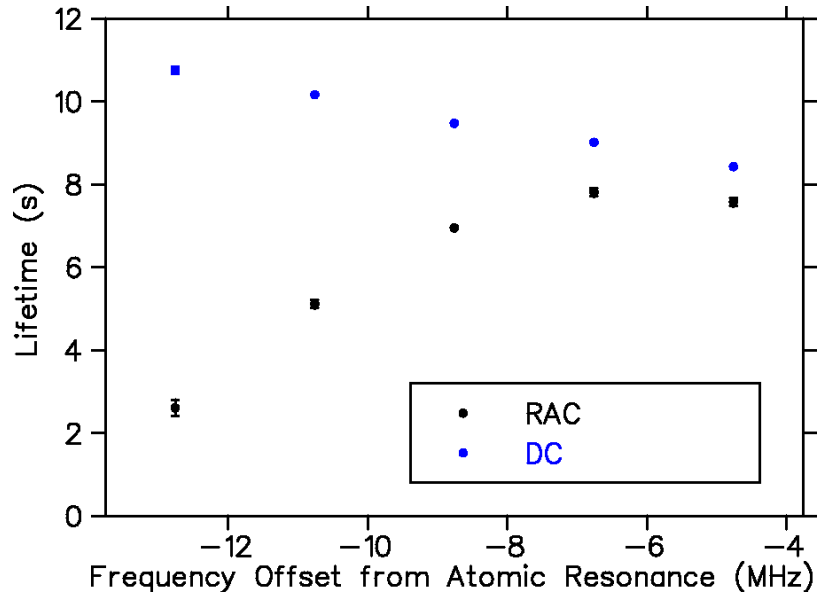

Figure 9: Plot of lifetime versus laser frequency (the frequency is given as an offset from the resonant transition $F = 2$ to $F = 3$ for potassium).
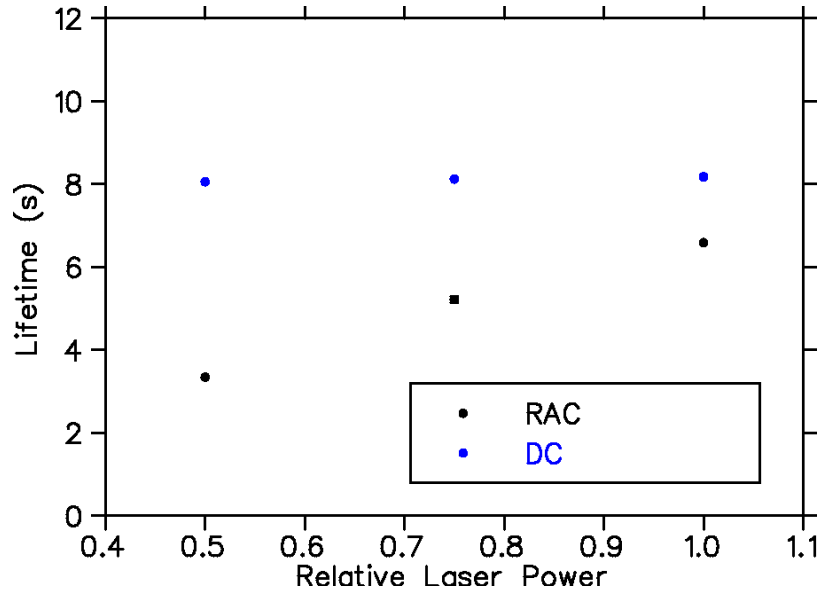
Figure 10: Plot of lifetime versus laser power. The laser power is given relative to the maximum laser power (normalized to 1).
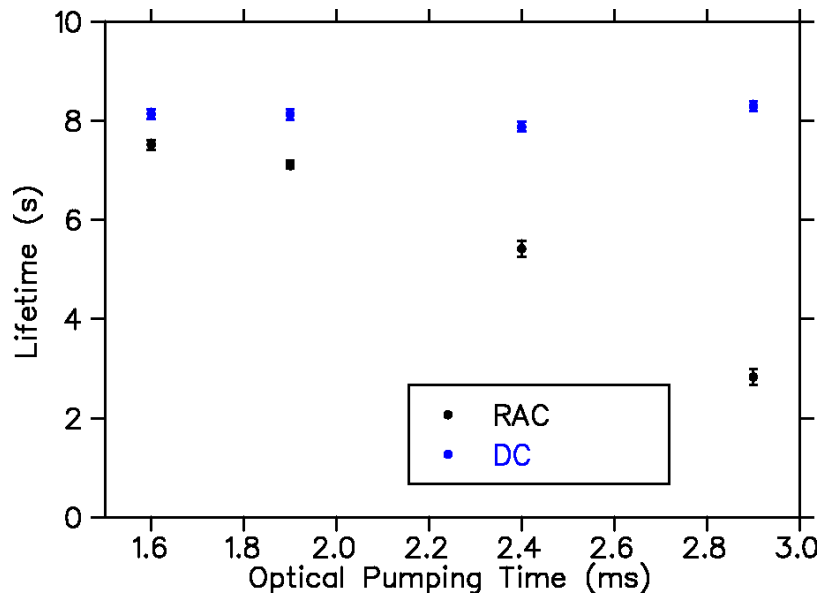

Figure 11: Plot of lifetime versus optical pumping time. Note that the DC values are plotted too, though "optical pumping time" is meaningless in the context of a DC trap. The values are plotted purely to show that the trap performance is not systematically changing for measurements made at different times.

3.3 Cloud Shape Measurements

3.3.1 What did we want to measure?

The IR camera is capable of taking exposures of around 90 μs at minimum. Since the trapping time for the RAC-MOT waveform is 2.9 ms, there exists the possibility of taking pictures during a single period of the RAC-MOT waveform and watching the cloud coalesce. We wanted to see how the width and centroid of the cloud changed over the course of a single trapping period.

We also took longer exposures to see if there was any variation in the cloud width and centroid over the full 480 ms cycle of the RAC-MOT (ideally, the cloud parameters should stay fairly constant). These longer exposures average the light collected from several periods of the RAC-MOT waveform.

3.3.2 How did we measure it?

Since the Pi simply waits after it sends the A0 signal telling the RAC-MOT to run a cycle, the Pi can be programmed to take a picture at a certain time in the cycle.

For the longer exposures, we used an exposure time of 25 ms — this averages the light over approximately 5 periods of the RAC-MOT waveform. We took 12 photos over a cycle of the RAC-MOT, each one at a different time in the cycle. To accomplish this, we made use of the variable **n** in the Pi code I mentioned previously (recall that **n** increments by one each time a cycle is complete, until it reaches **Max_n**, at which point it is reset to 1). We did not take all 12 photos within the same cycle; rather, we set **Max_n** to 12 and every time a new cycle was initiated, the Pi was made to delay by $35 \times$ **n** ms and then take a picture.

For the short exposures meant to capture the cloud formation during a single period of the RAC-MOT waveform, we used an exposure time of 0.5 ms. We took 6 photos over a period (2.9 ms), each one at a different time in the period. We used a similar strategy to the previous measurement, except that to ensure the timing was just right, we had the Pi monitor the BK TTL. After the first 5 periods had passed, the Pi would wait for the rising edge of the BK TTL, indicating the next period was about to begin. The Pi was made to wait $0.5 \times$ **n** ms following the rising edge and then take a picture.

We use a Physica code that takes vertical and horizontal projections of these pictures and fits a Gaussian function to both. The mean of the Gaussian gives the centroid value in that direction and the standard deviation of the Gaussian is taken as a measure of the width.

3.3.3 How can you do it yourself?

Relevant files:

Pi:StartSequenceRAC.sh Shell script which begins the Pi code and the camera code.

Pi:TrapSeq_cp3_RAC_cloudshape C++ code for the Pi to run for the long exposures for measuring the cloud shape over a cycle.

Pi: TrapSeq_cp3_RAC_cloudshape_1period C++ code for the Pi to run for the short exposures for measuring the cloud shape over a period.

Pi:params.txt System parameters.

trinatlt2:params.txt Camera parameters.

llawrence:cloudparams_gaussian.pcm Physica code to extract the cloud centroid and width from the pictures.

llawrence:cloudparamsplot_gaussian.pcm Physica code to plot the centroid and width as a function of time using the output data from **cloudparams.pcm**.

Recipe:

1. Open StartSequenceRAC.sh and change the code called on the Pi. The instruction may be found at the bottom of the file: change it to `sudo ./TrapSeq_cp3_RAC_cloudshape` or `sudo ./TrapSeq_cp3_RAC_cloudshape_1period`.

2. Open the params.txt file on the Pi and change **NumPulses** to 1. The variable **Max_n** for this code determines how many pictures are taken within a cycle or period. For pictures over the entire cycle, change **Max_n** to 12. For pictures over a single period, change **Max_n** to 6.

3. Open the params.txt file on trinatlt2 and change the exposure time, **expTimems** to 0.5.

4. Call `./StartSequenceRAC.sh` from the command line on the Pi to start the code.

5. The code will run indefinitely; type `q` and press Enter on trinatlt3 to stop the measurement.

6. Call `mvdata` on trinatlt2 and `yankdata [folder_name]` on the HP-Stream.

7. Enter the folder you just created with `yankdata`, where the copied pictures reside, and run Physica. Execute the code **cloudparams_gaussian.pcm** then **cloudparamsplot_gaussian.pcm**.

8. After the first run of **cloudparams_gaussian.pcm**, you only need to run **cloudparamsplot_gaussian.pcm** in the future to see the final results.

3.3.4   What were the results?

See Figures 12 and 13 for typical plots of the cloud parameters over one cycle and one period respectively. Each image contains four plots. The rows are organized by direction (x and y) and the columns by parameter (centroid and width). For example, the upper-right-hand plot shows the width in x. The white error bars indicate an uncertainty of the mean and the red error bars indicate the deviation of all samples taken at the same point in the cycle or period. See the Physica code for more details.
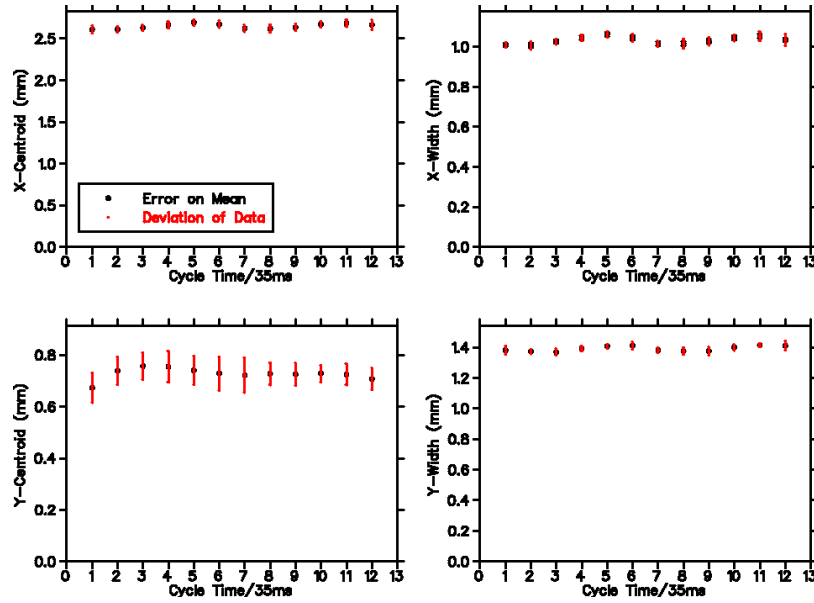
Figure 12: Cloud parameters as a function of cycle time using the usual half-sinusoid for the RAC-MOT waveform. The cloud parameters should be fairly constant over the duration of the cycle.
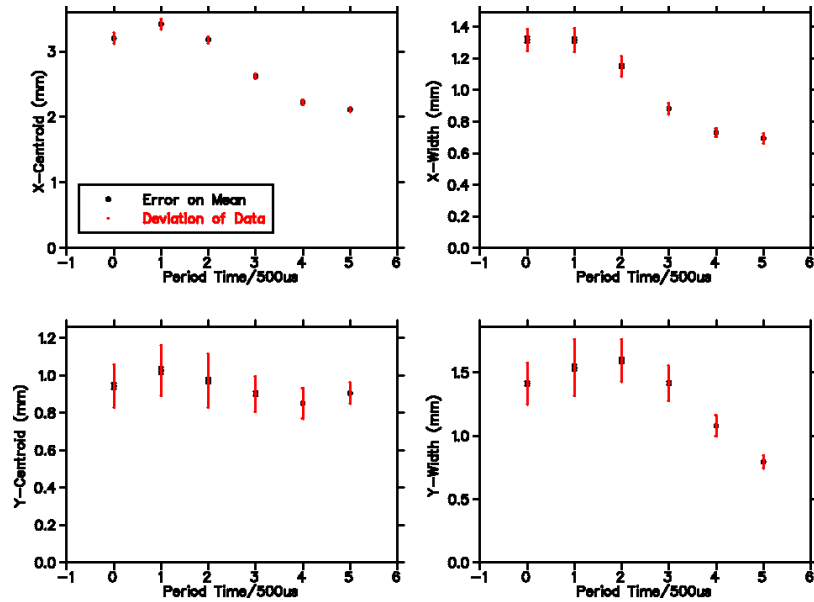

Figure 13: Cloud parameters as a function of period time using the usual half-sinusoid for the RAC-MOT waveform. The cloud should crunch down over the period (the width in x and y should decrease).
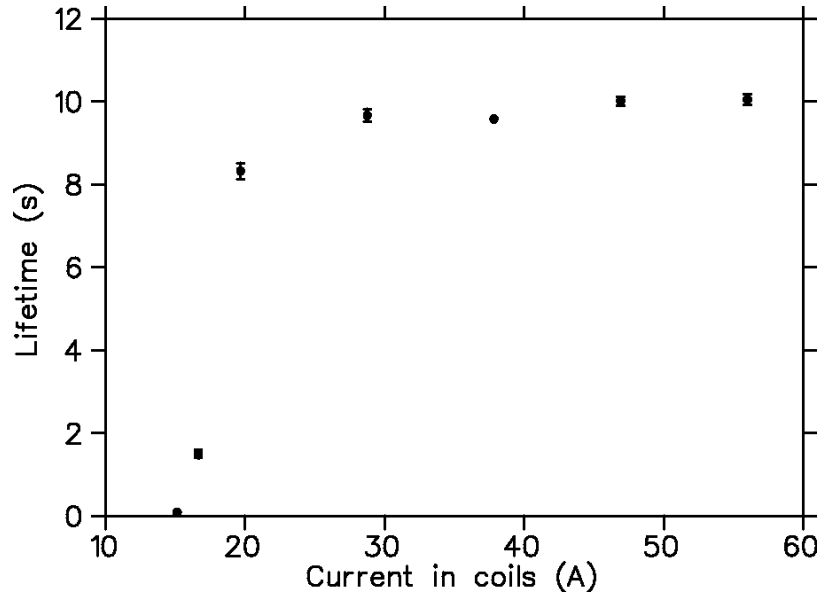
Figure 14: Plot of the lifetime of the trap versus the current in the main coils (and therefore the magnetic field strength) for a DC trap.

4   Magnetic Field Strength Considerations

The measurements described in the previous section, lifetime and cloud shape, serve as diagnostics for testing trap improvements. This section is one example of how we used these measurements to learn something about the behaviour of the trap.

The strength of the trap has an interesting relationship with the magnetic field strength. We took lifetime measurements of the DC trap for different settings of DC current in the main coils. The results are shown in Figure 14. There is a clear cut-off between 15 and 20 amps, below which the trap is very poor.

The RAC-MOT waveform during the trapping time spends a significant time at currents lower than this cut-off current. We attempted several variations on the usual waveform to make it spend less time at lower fields. We tried voltage waveforms which produce current waveforms in the Matsusadas as shown in Figure 15. I will refer to the black, red, and green waveforms as the sine, square, and sine-boost waveforms respectively (named for the input voltage waveforms, not the shape of the resultant currents).

Though the sine-boost waveform looks very ragged, the current does rise quickly above the cut-off current, and it was hypothesized that this might do a little better than the pure sine.

We took lifetime measurements for each of these waveforms. The results are shown in Table 2. We also took measurements of the cloud shape for each waveform; the plots are shown in Figures 16, 17, and 18.

Based on the lifetime and cloud shape measurements, we see that the sine-boost waveform does not perform as hoped: we do not get improved performance from the faster rise time (less time spend below the cut-off current). However, the square wave performs remarkably well, better than the other two waveforms in both lifetime and final cloud width (smaller width). The power
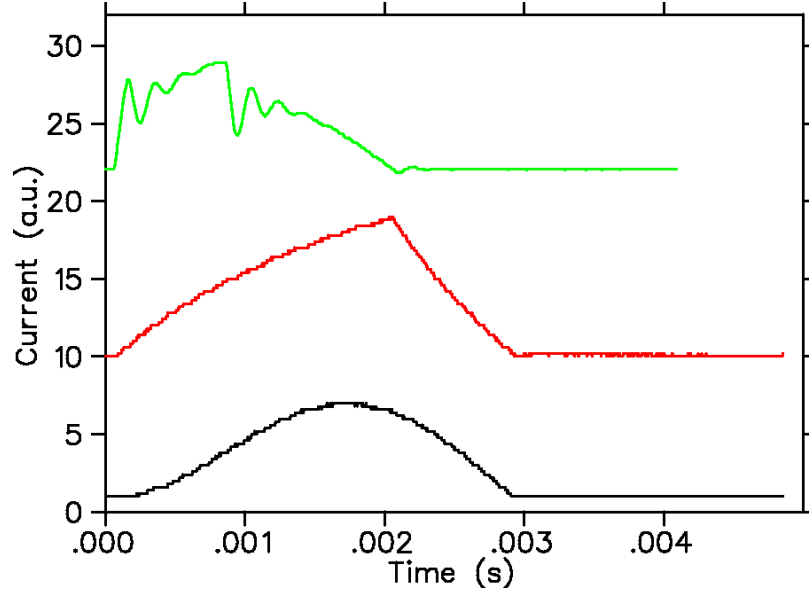
Figure 15: Current waveforms in Matsusada coils. The black waveform is the result of a sinusoidal voltage waveform. The red waveform is the result of a square voltage waveform. The green waveform is the result of the sum of a sinusoidal waveform and a boost pulse at the start of the sinusoid. Note that for the latter, the Matsusadas were in current output mode and the trapping time and off time were equal (2 ms).

Table 2: Best lifetime from a set of lifetime measurements for each of the aforementioned waveforms.

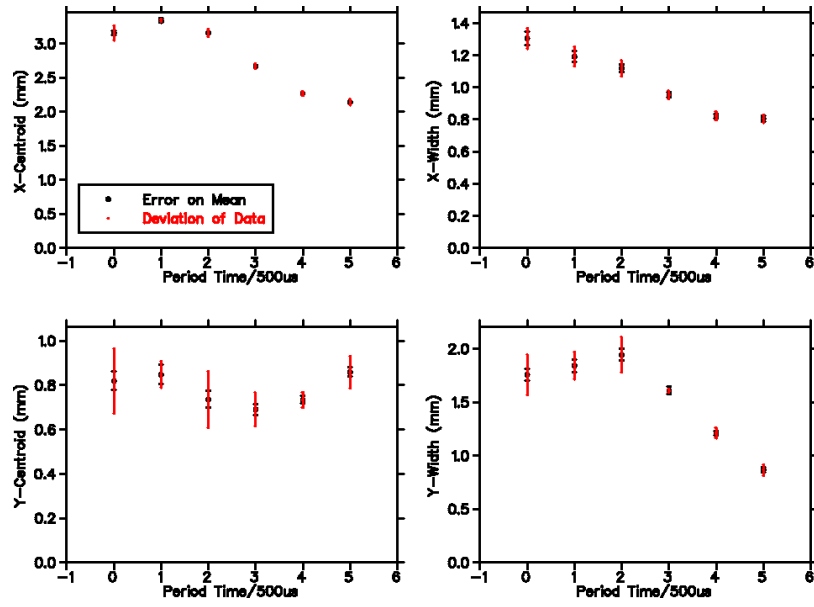| Waveform	Best Lifetime (s)
sine	9.6
square	11.3
sine-boost	7.4

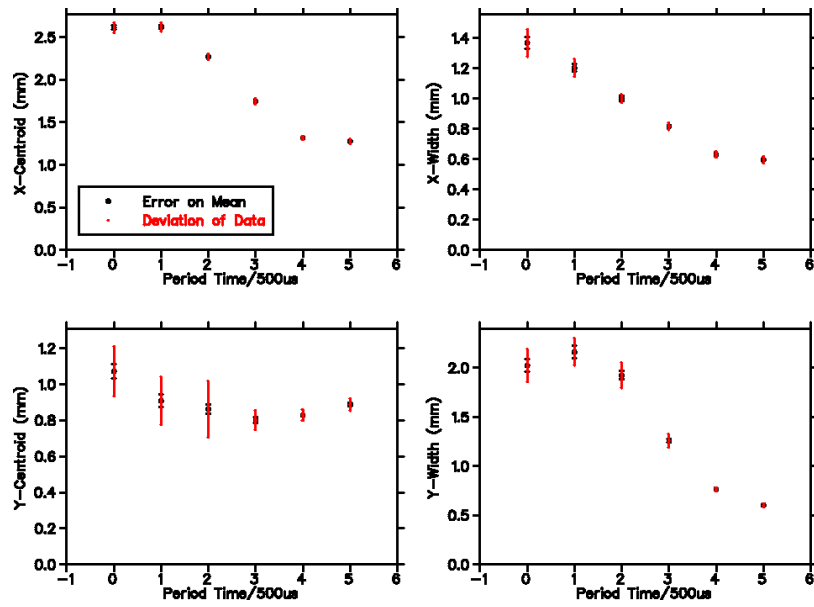Figure 16: Cloud shape measurements for the sine waveform.

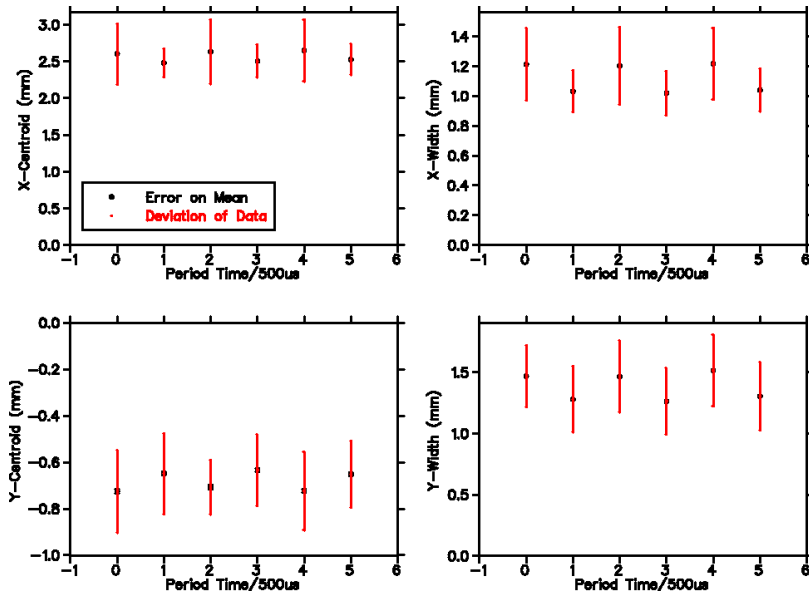Figure 17: Cloud shape measurements for the square waveform

Figure 18: Cloud shape measurements for the sine-boost waveform.

coupling for this waveform is unknown — if this waveform is to be used in the future, someone should calculate the relative power coupling using the technique described in the first section of this report.

A final note on the lifetimes and their relation to the magnetic field strength: is there a way one could calculate in advance the lifetime of the trap subject to a given waveform? The lifetime of the trap is usually on the order of *seconds*, while the current in the main coils varies over the course of *milliseconds*. It is a well-known fact from the theory of linear systems that if a first-order system is subject to a periodic input whose period is much less than the system's time constant, the system is unable to resolve the rapid fluctuations. The output will be indistinguishable from that produced if the input were replaced by a DC value equal to the waveform's average value.

Perhaps a similar principle applies to our MOT for a certain regime of waveform in the main coils. That is, the rapid fluctuations of the trapping forces (due to the rapid fluctuation of the current) is not resolvable by the trap loss mechanisms, and it is only the average trapping force that matters. If this is the case, we could predict the lifetime by calculating the average current for a given waveform and comparing this value to the plot of Figure 14.

We already have a source of data we can use to test this idea: we took lifetime measurements as a function of optical pumping time (see the previous section). We may calculate the average current of a function $f(t)$ with period $2T$ like so:

$$I_{avg} = \frac{1}{2T} \int_{-T}^{T} f(t) \, \mathrm{d}t. \tag{33}$$

The average value $A\sin(t)$ over half its period is $2A/\pi$. To obtain the average value for the RAC-MOT waveform, one must multiply by the duty cycle: the ratio of the trapping time to the entire period (trapping time plus optical pumping time). Let the trapping time be t_{trap} and the optical
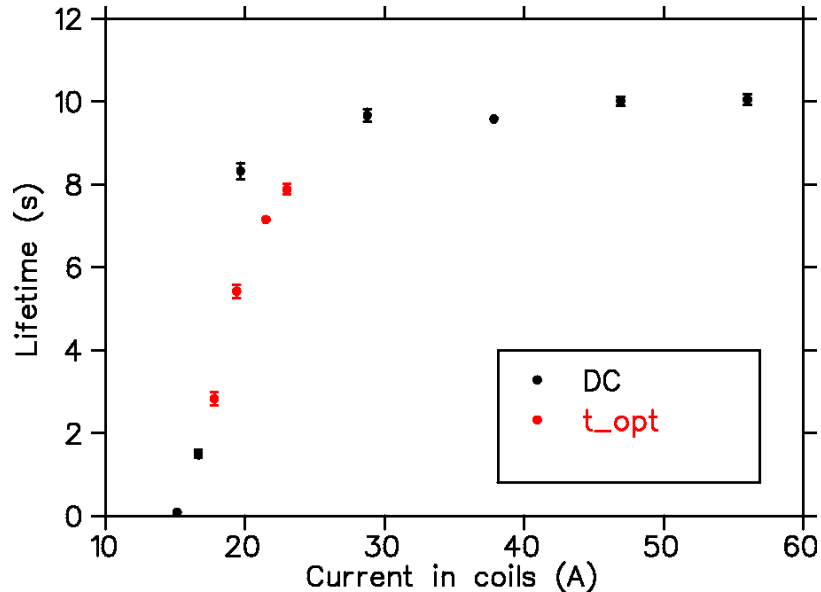
28

Figure 19: A plot of the lifetime versus average current. In black are the DC measurements and in red are the equivalent average currents calculated for the different optical pumping time measurements.

pumping time be t_{opt}. The duty cycle is:

$$D = \frac{t_{trap}}{t_{trap} + t_{opt}}. \tag{34}$$

The average current is:

$$I_{avg} = \frac{2A}{\pi} \frac{t_{trap}}{t_{trap} + t_{opt}} \tag{35}$$

As mentioned previously in this report, the amplitude is 56 amps and the trapping time is 2.9 ms. The optical pumping times measured are shown in Figure 11. From these values, we calculate the average current and plot the measured lifetime (from Figure 11) against the average current, along with the DC measurements of Figure 14. The results are shown in Figure 19.

There appears to be agreement here, given the optical pumping time measurements were taken more than a month before the DC field measurements and any number of things may have changed in the interim. More testing would have to be done, but if this relationship holds, it would provide a quick method of evaluating the efficacy of waveforms for trapping purposes.

This section concludes the description of the work done on the RAC-MOT during the summer of 2016. The remaining sections are a collection of miscellaneous ideas I had over the summer related to the TRINAT experiment.

5   Killing the Eddy Currents

While the RAC-MOT waveform helps with the heating problem in the chamber, there still remains the question of how to minimize the eddy currents during the optical pumping time (these eddy
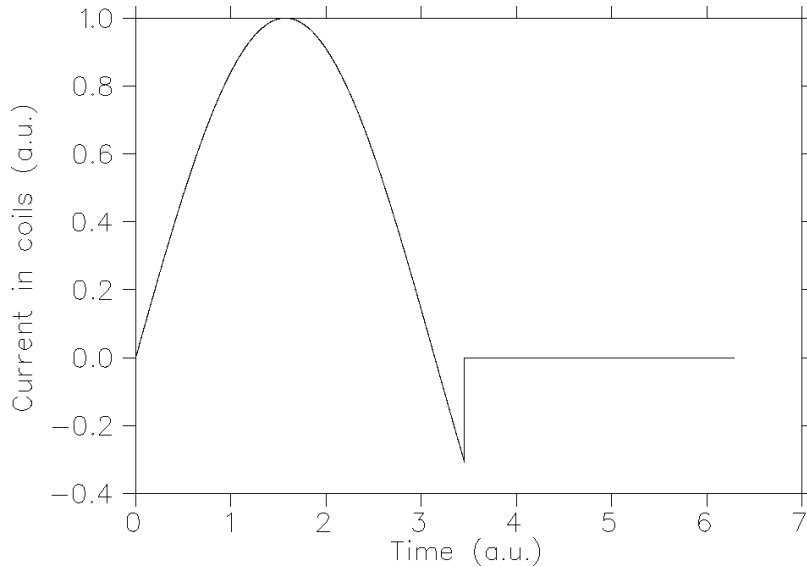
Figure 20: Ideal current waveform in main coils to kill the eddy currents.

currents lead to a magnetic field which spoils the polarization of the atoms). One proposal is to allow the current in the main coils a brief negative overshoot before turning off the Matsusadas — see Figure 20.

One can understand the idea intuitively like so: roughly, the eddy currents are proportional to the derivative of the current waveform in the main coils. While the sinusoid in the main coils has a negative slope, the eddy currents will be induced in one direction. If the sinusoid is brought to zero from its positive value (no overshoot), the eddy currents will decay exponentially from whatever value they had before the shut-off of the coil current. However, if the main coil current is allowed to overshoot negatively and then shut off, the derivative during the shut-off has a positive slope, inducing eddy currents in the *opposite* direction and cancelling the existing eddy currents. One could imagine tuning the timing of this negative overshoot in the lab to minimize the residual magnetic field from the eddy currents.

Unfortunately, there is a practical problem: the current in the main coils will itself decay exponentially because of the inductance of the coils. This current would also produce a non-dipole field (obviously) during the optical pumping time, spoiling the polarization. The way the electronics is set up now, if you shut off the driving waveform to the Matsusadas after a negative overshoot, the current in main coils would not look like that shown in Figure 20, but more like that shown in Figure 21.

The main coils each form an LR circuit which limits the maximum rate of change of the current:

$$L\frac{di}{dt} = -iR \tag{36}$$

The solution to this differential equation yields the exponentially decaying current seen in Figure 68:

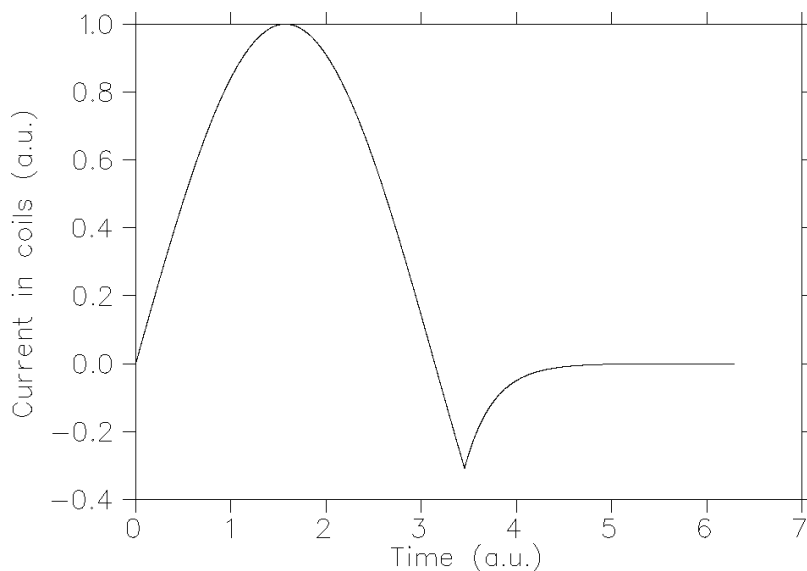$$i(t) = i(0)\exp\left(\frac{-tR}{L}\right) \tag{37}$$

30

Figure 21: The current waveform in the main coils upon shut-off decays exponentially because of the inductance of the coils.

So how can we improve this situation and hopefully obtain a current closer in character to that shown in Figure 20?

Notice what equation 36 says: since di/dt is proportional to the voltage across the coils, the maximum rate of change of the current is the product of the initial current $i(0)$ and the resistance R through which that current flows (generating a voltage). The voltage across the coils decreases as the current decays, which lessens di/dt over time, and leads to the leisurely exponential decay. Suppose we could pin the voltage across the coils at a given value? Then di/dt would be *constant*, and the current would ramp down linearly. If this pinned voltage is equal to or greater than $i(0)R$, this circuit would lead to a lower current at every point in time as compared to the LR circuit.

This idea comes from *The Art of Electronics*, page 53, in a discussion about inductors: "For fastest decay with a given maximum voltage, a zener could be used instead, giving a ramp-down of current rather than an exponential decay" [5]. A zener diode is a common component often used to clamp a voltage to a maximum value. The book goes on to mention the use of a metal-oxide varistor (MOV), which essentially acts as a bi-directional zener. Approximately, a MOV will conduct no current below a certain threshold voltage (I'll call it the clamping voltage); above this, the MOV will hold the voltage at its terminals at the clamping voltage.

In the context of this experiment, consider Figure 22. Assume we are running the RAC-MOT waveform, and let us consider what happens during the trapping time and the optical pumping time.

During the trapping portion of the cycle, the Matsusadas output a voltage waveform. The BK TTL is high, so the relay switch is closed. Suppose the MOV is selected such that its clamping voltage, V_{clamp} is above the maximum voltage output of the Matsusadas. This should not be a problem, as MOVs are available at "voltage ratings from 10 to 1000 volts and can handle transient currents up to thousands of amperes" [5]. During the trapping time, then, this circuit acts identically to the original circuit; one can ignore the relay and the MOV.
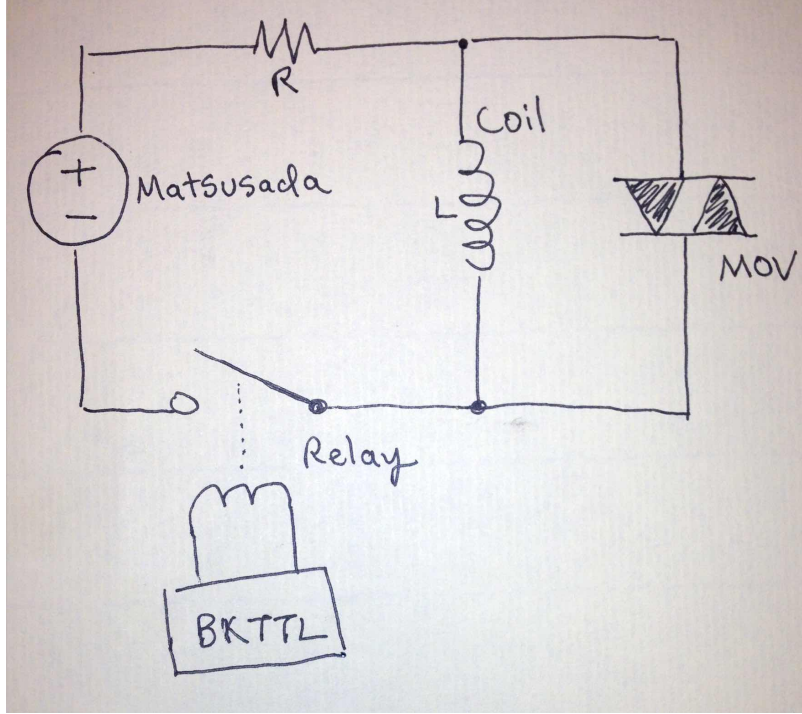
31

Figure 22: Sketch of circuit for fast shut-off of current in coils.

Now, when the optical pumping time begins, the BK TTL falls low, and the relay switch is opened. If there is still current flowing through the coils when this happens (i.e. if we allow the current a negative overshoot just before the optical pumping time), then the voltage across the coils will build and build as the the inductor looks for a way of pumping out the current. When this voltage exceeds V_{clamp}, the MOV "turns on" and the current may flow through this path, with the voltage across the coils clamped at V_{clamp}. The current in the coils will plunge to zero linearly, decaying to a negligible value faster than if the coils were made to pump out their current through the resistor.

To see this, consider the differential equation governing this process.

$$L\frac{di}{dt} = -V_{clamp} \tag{38}$$

The solution to this differential equation is:

$$i(t) = i(0) - \frac{V_{clamp}}{L}t \tag{39}$$

until $i(t)$ reaches zero. So long as $V_{clamp} > i(0)R$, $i(t)$ will be smaller at every point in time than if the inductor pumped out its current through the resistor.

I propose using such a circuit would allow the current waveform to resemble that shown in Figure 20 more closely, and the eddy currents could be killed before the optical pumping time begins as described previously. Note that the MOV would have to be chosen such that V_{clamp} is higher than the maximum voltage output of the Matsusadas, but lower than the sparking voltage

32

of the relay; otherwise, the current will jump across the open relay switch. You would want it as high as possible, though, for the fastest ramp-down of the current.

6 Ballistic Cloud Expansion and the Boltzmann Equation

In section 4.6.4 of Dan Melconian's doctoral thesis, he considers the ballistic expansion of the atom cloud when the trapping forces are turned off, and determines an expression for the relationship between the width of the atom cloud as a function of time and the temperature of the atoms [6]. Assuming the cloud begins as a 3-dimensional Gaussian, he proposes that the cloud will expand as a 3-dimensional Gaussian with a variance given by:

$$\boldsymbol{\sigma}^2(t) = \boldsymbol{\sigma}_0^2 + \frac{1}{2}v_0^2 t^2 \tag{40}$$

where $\boldsymbol{\sigma}^2(t) = (\sigma_x^2(t), \sigma_y^2(t), \sigma_z^2(t))$ is a vector of the variances in x, y, and z; $\boldsymbol{\sigma}_0^2$ is the initial variance vector; and v_0 is the most probable velocity of atoms in the gas. For a gas in thermal equilibrium, the most probable velocity is related to the temperature like so:

$$v_0^2 = \frac{2k_B T}{m} \tag{41}$$

where m is the mass of the atoms.

Dr. Melconian shows equation 40 to be reasonable using Monte Carlo simulations of cloud expansion for a given temperature. Claire Preston used his results to measure the temperature of the atom cloud by taking short-exposure pictures of the cloud after the trapping forces are turned off and fitting Gaussians to the x and y projections of the image. The temperature of the cloud is related to the variance of these projections, σ^2, as a function of time like so:

$$T = \frac{m}{k_B}\frac{d(\sigma^2)}{d(t^2)} \tag{42}$$

which may be easily shown from the previous equations.

I thought it would be interesting to try to derive equations 40 and 42 from basic theory to corroborate the simulation results. Derivation from basic theory might also give more insight into the physics. The following is what I came up with.

6.1 The Boltzmann Equation

We shall describe an ensemble of classical particles (all of mass m) by a single density function $f(\mathbf{x}; \mathbf{p}; t)$ defined over the 6-dimensional phase space of the particles — 3 dimensions for the position coordinates, 3 for the momentum components — at each point in time. Suppose we select a particle at random. The probability that our selected particle lies between positions $\mathbf{x} = (x, y, z)$ and $(x + dx, y + dy, z + dz)$ with momentum between $\mathbf{p} = (p_x, p_y, p_z)$ and $(p_x + dp_x, p_y + dp_y, p_z + dp_z)$ at time t is:

$$f(x, y, z; p_x, p_y, p_z; t)\mathrm{d}x\mathrm{d}y\mathrm{d}z\mathrm{d}p_x\mathrm{d}p_y\mathrm{d}p_z \tag{43}$$

The Boltzmann equation describes the evolution in time of this density function when the particles are subjected to an external force field K (this could be an applied electric field, for example). Here is a sketch of the derivation for one spatial dimension with $f = f(x; p_x; t)$. The underlying ideas are quite simple, making the resultant equation all the more elegant.

Consider the particles which exist at a certain point in phase space, with position x and momenta p_x. To where in the phase space do the particles move after a time δt? Assume there are no collisions (the particles do not interact).

Because of their momenta, the position of the particles will be changed to:

$$x + \frac{p_x}{m}\delta t \tag{44}$$

(using the Newtonian definition of momentum $p_x = mv_x$). Because of the external force field, the momenta of the particles will be changed to:

$$p_x + K\delta t. \tag{45}$$

In the context of our density function:

$$f(x; p_x; t) = f(x + \frac{p_x}{m}\delta t; p_x + K\delta t; t + \delta t) \tag{46}$$

The density which at time t existed at $(x; p_x)$ may be found at time $t + \delta t$ at $(x + \frac{p_x}{m}\delta t; p_x + K\delta t)$. We expand the right-hand side in a Taylor series:

$$f(x; p_x; t) = f(x; p_x; t) + \frac{\partial f}{\partial x}\frac{p_x}{m}\delta t + \ldots + \frac{\partial f}{\partial p_x}K\delta t + \ldots + \frac{\partial f}{\partial t}\delta t + \ldots \tag{47}$$

The ellipses denote the higher-order terms in the series (all containing terms of $(\delta t)^2$ or higher). We cancel the term $f(x; p_x; t)$ on both sides, divide by δt, and take the limit as $\delta t \to 0$. The result is the collision-less Boltzmann equation in one dimension.

$$\frac{\partial f}{\partial t} + \frac{p_x}{m}\frac{\partial f}{\partial x} + K\frac{\partial f}{\partial p_x} = 0 \tag{48}$$

This generalizes naturally to 3 spatial dimensions:

$$\frac{\partial f}{\partial t} + \frac{\mathbf{p}}{m} \cdot \nabla_{\mathbf{x}}f + \mathbf{K} \cdot \nabla_{\mathbf{p}}f = 0 \tag{49}$$

The effect of collisions is to add a forcing term to the right-hand side of the equation. We are only interested in ballistic expansion here (no collisions), so I will consider this term to be zero[4].

---

[4]An interesting aside: you can show that this term is zero for point particles which interact under elastic collisions, since the only effect of a collision is to swap the velocities. Swapping velocities is equivalent to the particles passing right through each other for our density function f, which only measures the number of particles at a given point in phase space. So a group of elastically-colliding point particles is the same as a group of non-interacting point particles.

6.2 Ballistic Expansion

We now apply the Boltzmann equation and density function to describe the ballistic expansion of a cloud of particles with no external force field present. We assume the particles have a known initial density. Since we assume no external forces, the last term in equation 49 drops out. We are left with:

$$\frac{\partial f}{\partial t} + \frac{\mathbf{p}}{m} \cdot \nabla_{\mathbf{x}} f = 0. \tag{50}$$

Since we do not need to consider forces, let us use velocity instead of momentum so that $f = f(\mathbf{x}; \mathbf{v}; t)$, and the preceding equation becomes:

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{x}} f = 0. \tag{51}$$

Let the expansion begin at time $t = 0$. Suppose we have initial conditions $f(\mathbf{x}; \mathbf{v}; 0) = \phi(\mathbf{x}; \mathbf{v})$. Since there are no forces and no collisions to change a particle's velocity, the density will not change in the velocity coordinates over time. Additionally, a particle which is located at $\mathbf{x}_0$ at $t = 0$ will be located at $\mathbf{x}_0 + \mathbf{v}t$ at time t. The general solution to equation 51 is thus:

$$f(\mathbf{x}; \mathbf{v}; t) = \phi(\mathbf{x} - \mathbf{v}t; \mathbf{v}). \tag{52}$$

This satisfies the initial conditions and the differential equation, since

$$\frac{\partial f}{\partial t} = \frac{\partial}{\partial t} \phi(\mathbf{x} - \mathbf{v}t; \mathbf{v}) = -\mathbf{v} \cdot \nabla_{\mathbf{x}} f \tag{53}$$

and

$$f(\mathbf{x}; \mathbf{v}; 0) = \phi(\mathbf{x}; \mathbf{v}). \tag{54}$$

In the lab, we are often not interested in f per se, since we generally measure position density only and not velocity density. That being the case, we wish to determine the *marginal* density of f with respect to the position variables of $\mathbf{x}$ — I will call this marginal density $g(\mathbf{x}; t)$. This is done by integrating over the velocity variables. It turns out this integration takes on a simple and illuminating form if we use Cartesian coordinates and make the assumption that the initial density is separable in every variable. That is, we assume we may write the initial density like so:

$$\phi(\mathbf{x}; \mathbf{v}; t) = X(x)Y(y)Z(z)V_x(v_x)V_y(v_y)V_z(v_z) \tag{55}$$

(you will see that this is a good assumption for our purposes). Let this be the case, and integrate out the velocity variables from the solution of f to obtain g (the limits of integration are $(-\infty, \infty)$).

$$
\begin{aligned}
g(x, y, z; t) &= \int \phi(\mathbf{x} - \mathbf{v}t; \mathbf{v}) \, \mathrm{d}^3 v & (56) \\
&= \int \int \int X(x - v_x t)Y(y - v_y t)Z(z - v_z t)V_x(v_x)V_y(v_y)V_z(v_z) \, \mathrm{d}v_x \mathrm{d}v_y \mathrm{d}v_z & (57) \\
&= \int X(x - v_x t)V_x(v_x) \, \mathrm{d}v_x \int Y(y - v_y t)V_y(v_y) \, \mathrm{d}v_y \int Z(z - v_z t)V_z(v_z) \, \mathrm{d}v_z & (58)
\end{aligned}
$$

Each of these separated integrals has the same form, so consider just the one in x. Make the substitution $u = v_x t$.

$$\int X(x - v_x t) V_x(v_x) \, \mathrm{d}v_x = \frac{1}{t} \int X(x - u) V_x \left(\frac{u}{t} \right) \, \mathrm{d}u \tag{59}$$

The right-hand side of this equation is a convolution integral! Recall that the convolution of two functions $f(x)$ and $g(x)$ is defined by:

$$f(x) * g(x) = \int f(x - u) g(u) \, \mathrm{d}u \tag{60}$$

The integral in equation 59 is simply:

$$\frac{1}{t} X(x) * V_x \left(\frac{x}{t} \right) \tag{61}$$

And therefore, the marginal density $g(x, y, z; t)$ is given by:

$$g(x, y, z; t) = \frac{1}{t^3} \left(X(x) * V_x \left(\frac{x}{t} \right) \right) \left(Y(y) * V_y \left(\frac{y}{t} \right) \right) \left(Z(z) * V_z \left(\frac{z}{t} \right) \right) \tag{62}$$

6.3 Expansion of the Atom Cloud

With equation 62, we are suitably equipped to tackle the original problem: to determine the relation between the width of the atom cloud as a function of time and the temperature of the atoms. We just need to know the initial conditions and we are set to calculate.

We shall assume the initial velocity distribution to be the Maxwell-Boltzmann distribution, the classical velocity distribution for an ensemble of particles in thermal equilibrium. Right away, it must be said that this might be problematic, since the Maxwell-Boltzmann distribution is derived assuming a gas confined in a box, with no other forces present. The atoms in the MOT experience a harmonic trapping potential. Dr. Melconian mentions the same thing in his thesis; however, it appears to be a common procedure to simply assume a Maxwell-Boltzmann distribution, as another paper I read on the subject does the same thing [7].

The Maxwell-Boltzmann distribution gives the probability of a particle having a velocity between (v_x, v_y, v_z) and $(v_x + \mathrm{d}v_x, v_y + \mathrm{d}v_y, v_z + \mathrm{d}v_z)$ as:

$$P_v(v_x, v_y, v_z) \mathrm{d}v_x \mathrm{d}v_y \mathrm{d}v_z = \left(\frac{1}{2\pi a^2} \right)^{3/2} \exp\left(-\frac{v_x^2 + v_y^2 + v_z^2}{2a^2} \right) \mathrm{d}v_x \mathrm{d}v_y \mathrm{d}v_z \tag{63}$$

where

$$a^2 = \frac{kT}{m} \tag{64}$$

(note that this notation is different from that used in Dr. Melconian's thesis).

For the initial position density, we have observed that it is well-described by a three-dimensional Gaussian. Let the principal axes lie upon x, y, and z.

$$P_x(x, y, z) \mathrm{d}x \mathrm{d}y \mathrm{d}z = \left(\frac{1}{(2\pi)^{3/2} \sigma_x \sigma_y \sigma_z} \right) \exp\left(-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2} - \frac{z^2}{2\sigma_z^2} \right) \mathrm{d}x \mathrm{d}y \mathrm{d}z \tag{65}$$

Our initial condition is given by the product of P_v and P_x — a 6-dimensional Gaussian expressed along its principal axes. Notice that this initial condition is separable in all variables, as required by our solution to the Boltzmann equation under ballistic expansion, equation 62. Let us examine just one of the convolution integrals of equation 62 using this initial condition, the one for (x, v_x), for instance.

$$\frac{1}{t}X(x) * V_x\left(\frac{x}{t}\right) = \frac{1}{\sqrt{2\pi}\sigma_x}\exp\left(\frac{-x^2}{2\sigma_x^2}\right) * \frac{1}{\sqrt{2\pi}at}\exp\left(\frac{-x^2}{2a^2t^2}\right) \tag{66}$$

We see we have the convolution of two Gaussians, one with a variance σ_x^2 and one with a variance a^2t^2! It is well-known that the convolution of two Gaussians with variances σ_1^2 and σ_2^2 is another Gaussian whose variance is equal to the sum $\sigma_1^2 + \sigma_2^2$ (see, for instance, [3]). Therefore, the position density $g(x, y, z; t)$ is given by:

$$g(x, y, z; t) = \frac{\exp\left(-\frac{x^2}{2(\sigma_x^2 + a^2t^2)} - \frac{y^2}{2(\sigma_y^2 + a^2t^2)} - \frac{z^2}{2(\sigma_z^2 + a^2t^2)}\right)}{(2\pi)^{3/2}\sqrt{(\sigma_x^2 + a^2t^2)(\sigma_y^2 + a^2t^2)(\sigma_z^2 + a^2t^2)}} \tag{67}$$

The cloud does indeed expand as a 3-dimensional Gaussian, with the variance vector $\sigma^2(t)$ as a function of time equal to:

$$\boldsymbol{\sigma}^2(t) = \boldsymbol{\sigma}_0^2 + a^2t^2 = \boldsymbol{\sigma}_0^2 + \frac{k_B T}{m}t^2 \tag{68}$$

Therefore, we may choose any of the principal axes of the Gaussian (x,y,or z) and track the expansion of the cloud along this axis as a function of time (as Claire Preston did with the x and y projections of the IR images). From the preceding equation, the measured variance along the axis, $\sigma^2(t)$, is related to the temperature like so:

$$T = \frac{m}{k_B}\frac{d(\sigma^2)}{d(t^2)} \tag{69}$$

as we set out to prove.

6.4   Some Comments

Although this derivation from basic theory yields the same results as the simulations, it also gives a little more insight into the problem. I would like to use this section to discuss some of these insights.

We see that it is something of an accident that the cloud expands as a 3-dimensional Gaussian: this is not a general property. This is only the case because we have assumed the initial position density is well-described by a 3-dimensional Gaussian. The convolution of the Gaussian velocity distribution with the Gaussian position distribution yields another Gaussian whose variance changes in time. If the initial position density were some other function, the position density of the atoms upon expansion would not be described by a Gaussian.

What can we say about the expanding cloud if its initial position density is not Gaussian? It is a general theorem that variances add under convolution [3]. What I mean is this: define the mean μ and variance σ^2 of a normalized function $f(x)$ as

$$\mu = \int xf(x)\,dx \tag{70}$$

and

$$\sigma^2 = \int (x - \mu)^2 f(x) \, \mathrm{d}x. \tag{71}$$

Let the variance of $f(x)$ be σ_f^2 and that of $g(x)$ be σ_g^2. If $h(x)$ is the convolution of $f(x)$ with $g(x)$, then the variance of $h(x)$, σ_h^2, is given by:

$$\sigma_h^2 = \sigma_f^2 + \sigma_g^2. \tag{72}$$

Even if the initial position density (and thus the density under expansion) is not well-described by a Gaussian, so long as that function is separable along the position coordinates, we may still relate the variance of this function — as defined by equation 71 — along these position coordinates to the temperature using equations 68 and 69. These expressions are more general than equation 67, and might be useful for future experiments. One could calculate the variance approximately from a discrete data set (e.g. the projections of the IR camera images) by using

$$\mu = \sum_{i=1}^{n} x_i p_i \tag{73}$$

and

$$\sigma^2 = \sum_{i=1}^{n} (x_i - \mu)^2 p_i \tag{74}$$

where p_i is the normalized pixel intensity of pixel x_i (normalized in the sense that the sum of all pixel intensities is one). This would circumvent the need to fit a function to the data, handy if there is no simple function which fits well.

References

[1] M. H. Anderson, W. Petrich, J. R. Ensher, and E. A. Cornell. Reduction of light-assisted collisional loss rate from a low-pressure vapor-cell trap. *Phys. Rev. A*, 50:R3597–R3600, Nov 1994.

[2] Melissa Anholm. *Characterizing the AC-MOT*. PhD thesis, University of British Columbia, 2011.

[3] Ronald Bracewell. *The Fourier Transform and Its Applications*. McGraw-Hill, 1999.

[4] Matthew Harvey and Andrew James Murray. Cold atom trap with zero residual magnetic field: The ac magneto-optical trap. *Phys. Rev. Lett.*, 101:173201, Oct 2008.

[5] Paul Horowitz and Winfield Hill. *The Art of Electronics*. Cambridge University Press, second edition, 1989.

[6] Dan Melconian. *Measurement of the neutrino asymmetry in the beta decay of laser-cooled, polarized 37K*. PhD thesis, Simon Fraser University, 2005.

[7] I. Yavin, M. Weel, A. Andreyuk, and A. Kumarakrishnan. A calculation of the time-of-flight distribution of trapped atoms. *Am. J. Phys.*, 2001.