LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note	**LIGO-T060096-01-Z**

Calibration of ALLEGRO for the stochastic search

Martin McHugh, and
Warren Johnson, Jon Hanson, Phil Miller, Damon Nettles, Andrew Weber, and Jason Weaver
Loyola of New Orleans, and Louisiana State University

Distribution of this document:

Stochastic Analysis Group

Draft

This is an internal working
note of the LIGO project

California Institute of Technology
LIGO Project, MS 18-34
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW17-161
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

WWW: http://www.ligo.caltech.edu/

CVS Id: ALLEGROCalibTechNote.tex,v 1.32 2006/09/27 13:44:03 mchugh Exp — Processed with LaTeX on 2006/09/28

Abstract

We describe the methods and procedures used to calibrate the ALLEGRO resonant detector during the S4 run, when we made a search for stochastic gravitational waves, in collaboration with the LIGO Livingston Observatory. First, the interaction of the gravitational field with the resonant detector is calculated. Then the complete antenna system is described. Then we describe the "calibrator", or force actuator, and the methods used to determine its actuation constant. This then allows an absolute determination of the transfer function (or response function) that connects the gravitational strain to the observable output. We then explain how this response function is used to back-transform the raw observational data to nominal gravitational strain. A comprehensive discussion of experimental detail is included, with many links to relevent log entries.

Contents

1 Introduction

The ALLEGRO resonant detector, operated by a group from Louisiana State University [1], is a two-ton aluminum cylinder coupled to a niobium secondary resonator. The secondary resonator is part of an inductive transducer [2] which is coupled to a DC SQUID. Raw data acquired from the detector thus reflects the high Q resonant mechanical response of the system. We must undo this transformation to get back the spectrum of the gravitational wave strain signal. Previous joint searches for gravitational waves involving ALLEGRO have focused on burst-like sources [3]. In these searches the raw detector output is filtered [4][5] for burst-like signals. The filtering produces an event list of 'triggers' – calibrated for some standard choice of burst waveform – which is compared with triggers generated in other detectors. No coherent method is applied to multiple detectors in these searches. In contrast, with a search for stochastic sources of gravitational waves [6], such as a primordial background [7], the coherence of the stochastic signal must be preserved in each detector's calibrated data. Such a search is being undertaken [8][9] by cross-correlating data from ALLEGRO and the LIGO Livingston interferometer [10]. Some details of the technique employed for this search are presented by Whelan *et al* [11].

In this note we describe the procedures and measurements used to calibrate the resonant detector. In logical order, we first determine the precise theoretical correspondence between: 1) a gravitational field $h(t)$ and 2) a force $F(t)$ applied to *one* end of the bar (Sec. 4.1). Then we measure the absolute "actuation" constant of our force generator, which is a "calibration" capacitor at one end of the bar (Sec. 5.2.2). This allows us to apply a known $F(t)$ to one end of the bar. Finally, we measure the absolute response function (or transfer function) that transforms force into the observable voltage. A key element of this is getting all of the timing and phase relationships right so that any signal can be faithfully recovered (Sec. 5.1).

All of the above are then used to make the crucial inversion, or back-transformation, from the recorded detector output to nominal gravitational strain $h(t)$.

2 The mechanics of a resonant detector

To first approximation, the mechanics of ALLEGRO are described in Fig. 1. The antenna proper is a right-circular cylinder, or "bar", and is represented in the schematic as the masses m_2 and m_3 and the spring k_2. They form a mechanical harmonic oscillator that is equivalent to the first longitudinal mode of vibration of the actual bar, with a frequency near 904 Hz. The mechanical resonator, or secondary resonator (m_4 and k_3), is shown attached to the right face of the bar; it is tuned to have nearly the same resonant frequency, and it acts as a mechanical transformer, converting small motions of the big masses (m_2 and m_3) into

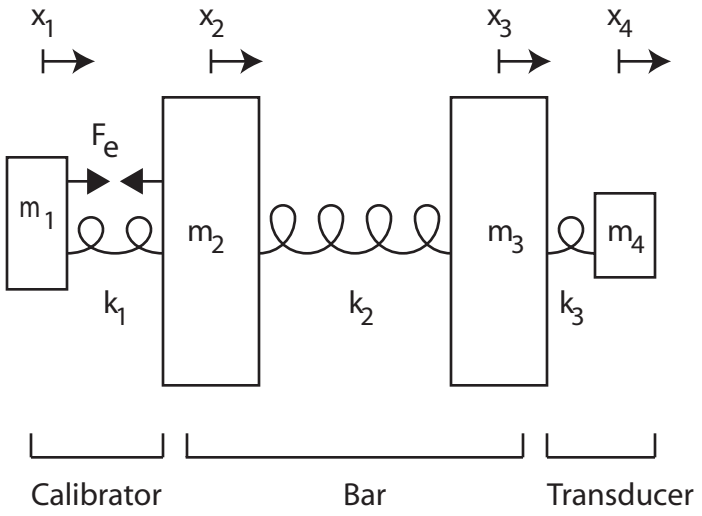

Figure 1: A schematic diagram of the mechanical parts of ALLEGRO, to first approximation

big motions of the smaller mass m_4. The inductive transducer and amplifier (the SQUID) produce an output voltage proportional to the differential displacement $x_4(t) - x_3(t)$.

The calibrator is a capacitor plate m_1 attached to the other face of the bar by a weak spring k_1. It faces a second capacitor plate (not shown), firmly attached to the end of the bar. AC and DC voltages are applied to those plates to produce equal but opposite electrostatic forces $F_e(t)$. Because the spring k_1 is very weak (the resonant frequency is 200 Hz), a force applied near 900 Hz on the reaction mass m_1 (inner capacitor plate) is almost entirely converted into inertial motion, and essentially none of it is communicated back to the face of the bar via the mechanical spring.

3 The complete detector

The path of a signal through the ALLEGRO detector is represented schematically in Fig. 2. Each block in the diagram represents a particular transfer function that forms a stage in the overall response of the detector. The first element shown is the gravitational field, represented by the strain $h(t)$. It produces any true gravitational part of the total force $F(t)$, conventionally assumed to be applied to one end of the bar. This force is simply proportional to the second time derivative of the strain. All end forces excite the longitudinal vibrational mode of the bar, which is strongly coupled to a secondary mechanical resonator mounted at the right end. Motion of this secondary resonator $x(t)$, relative to the bar face, drives a current $I(t)$ in a superconducting inductive transducer. A DC SQUID operating in a flux-locked loop acts as a flux-to-voltage transducer. Its output voltage (v_S) is then demodulated by a two channel lock-in amplifier. The lock-in reference oscillator is synchronized to a frequency source which is synchronized to a GPS receiver. The lock-in is used to heterodyne the output of the SQUID, demodulating from $\sim$ 900 Hz down to DC, thus reducing the required sampling rate and data volume. The value of this reference frequency is chosen to

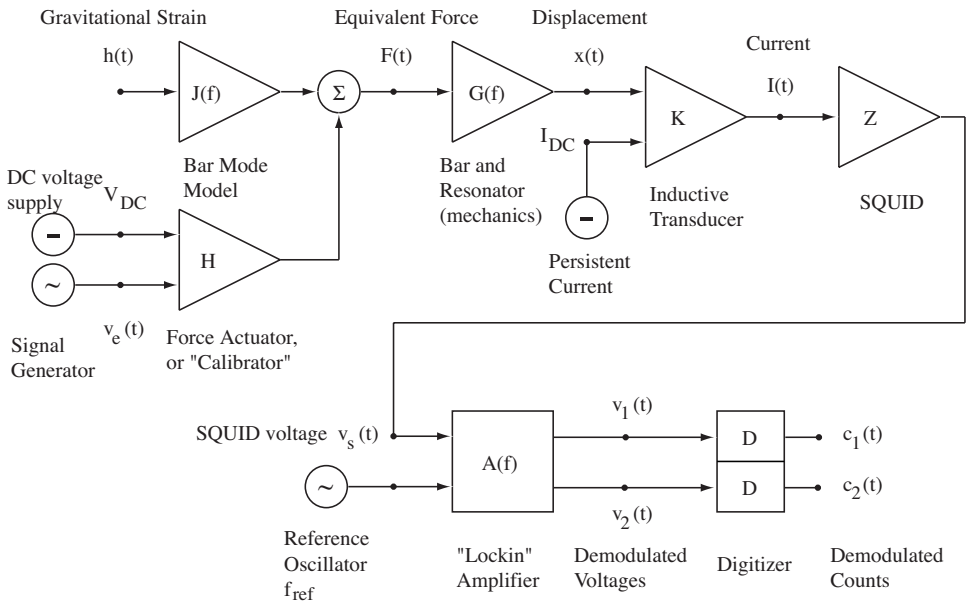

Figure 2: Diagram of the signal path for the complete detector. The blocks are components that can be represented by linear input-output relations. The dots represent inputs or outputs that may or may not be directly observable, but are useful for understanding the dynamics of the system. All of the mechanics are included in the block labeled G. Unlike interferometric gravitational wave detectors, feedback plays only a minor role.

be near the center of the sensitive band, about half-way between the normal mode frequencies of the bar-resonator pair. For the S4 data run, this value was fixed at 904Hz.

The in-phase and quadrature outputs of the lock-in (v_1, v_2) go to low-pass anti-aliasing filters and then to a data acquisition system which samples and stores the data (c_1, c_2) to a hard disk at a fixed rate of 250 samples/s.

3.1 Feedback

There is an undesired, but weak, direct feedback from the SQUID input to the mechanical resonator. Its effect is to "pull" the normal frequencies by a small amount (< 80 mHz), and to destabilize the resonances. To counteract this, another feedback path has been added from the SQUID output to the calibrator at the other end of the bar, to add some small negative force feedback ("electronic damping") to the modes. Both feedback couplings are weak, and their effects are adequately summarized by determining the experimental mode frequencies and decay times. We have thus omitted both feedback paths from the diagram and from further discussion.

4 Transfer functions

Here we explain the frequency domain transfer functions represented by the triangles and boxes in Fig. 2.

4.1 Strain to Force

The first of these is the strain to force relationship $J(f)$. The gravitational wave strain, having tensor components $\ddot{h}_{ij}$, induces a tidal force distributed throughout the bar with the force density at a coordinate x_j in a bar of mass density ρ given by

$$f_i^{\mathrm{GW}} = \frac{\rho}{2} \sum_j \ddot{h}_{ij} x_j \tag{4.1}$$

We represent the displacement of an element of an elastic body with u. This obeys a differential equation[12]

$$\ddot{u}_i = (\lambda + \mu)\nabla_i(\nabla \cdot \mathbf{u}) + \mu \nabla^2 u_i + f_i^{\mathrm{GW}} + f_i^X \tag{4.2}$$

Where λ and μ are the Lame form of the elastic moduli. The elastic body equation has solutions of the form

$$\mathbf{u} = \sum_m a_m(t)\Psi_m(\mathbf{x}) \tag{4.3}$$

We will only be concerned with the first longitudinal mode. The cylinder axis is denoted as x, with $x = 0$ at the center of the bar. The term $\Psi_1(x)$ is the mode shape, and for this first mode we have $\Psi_1(x) = \sin\frac{\pi x}{L}$ where L is the physical length of the bar. We have $\Psi_1(x = L/2) = 1$ so this means $a_1(t)$ is the amplitude of motion of the end of the bar. Our differential equation now becomes

$$\ddot{a}_1(t) + \omega_1^2 a_1(t) = \frac{1}{\rho N_1} \int (f^{\mathrm{GW}} + f^X) \cdot \Psi_1 d^3x \tag{4.4}$$

The term N_1 is a normalization constant ($N_1 \equiv \int \Psi_1 \cdot \Psi_1 d^3x$). This is a volume integral, and integrating over two of the three dimensions gives us the cross sectional area of the bar (A). So we have

$$N_1 = A \int_{-L/2}^{L/2} \sin^2 \frac{\pi x}{L} dx = \frac{AL}{2} = \frac{V}{2} \tag{4.5}$$

The normalization is equal to half of the volume of the bar. We now evaluate the right hand side of Eq. 4.4. We go back to our force density from the gravitational wave from Eq. 4.1. For the external force density f^X we choose a force that approximates the force applied by the calibrator. This is a force applied to one end of the bar at a single point

($f^X = F_c \cdot \delta_x(x - L/2) \cdot \delta_y(y) \cdot \delta_z(z)$). Our equation is now

$$\begin{aligned} \ddot{a}_1(t) + \omega_1^2 a_1(t) &= \frac{2}{\rho V} \int_{-L/2}^{L/2} (\frac{1}{2}\rho \ddot{h} x + F_c \cdot \delta_x(x - L/2) \cdot \delta_y(y) \cdot \delta_z(z)) \cdot \sin \frac{\pi x}{L} d^3x \\ &= \frac{2}{\rho V}[\frac{V \rho \ddot{h} L}{\pi^2} + F_c] \\ &= \frac{2\ddot{h}L}{\pi^2} + \frac{2F_c}{M} \end{aligned} \tag{4.6}$$

We have two terms on the right hand side, the first for the distributed tidal force due to the strain, the second for a direct force applied to a point on the end of the bar. It is this second force which we can then use with the mechanical model discussed in Section 2. We can replace the strain term with an effective force applied to the end of the bar with the following.

$$\ddot{h} = \frac{\pi^2 F_{\text{eff}}}{ML} \tag{4.7}$$

In the frequency domain this becomes

$$\begin{aligned} \tilde{h} &= \frac{-\pi^2 \tilde{F}_{\text{eff}}}{\omega^2 ML} \\ \tilde{h} &= \frac{-\tilde{F}_{\text{eff}}}{4f^2 ML} \end{aligned} \tag{4.8}$$

Inverting gives our transfer function

$$\begin{aligned} \tilde{F}_{\text{eff}} = -4f^2 ML\tilde{h} = J(f)\tilde{h} \\ J(f) = -4f^2 ML \end{aligned} \tag{4.9}$$

The mass M and length L of the aluminum bar are known parameters.

4.2 Force to displacement

Next we deal with the mechanical response of the system to this effective force. We use a model based on the mechanics discussed in Sec. 2. We have a pair of damped coupled oscillators.

$$\begin{aligned} \tilde{x}(f) &= G(f)\tilde{F}(f) \\ \tilde{x}(f) &= \alpha \left(\frac{1}{\left(f_p^2 - f^2 + \frac{if_p f}{Q_p}\right)} - \frac{1}{\left(f_m^2 - f^2 + \frac{if_m f}{Q_m}\right)} \right) \tilde{F}(f) \\ G(f) &= \alpha \left(\frac{1}{\left(f_p^2 - f^2 + \frac{if_p f}{Q_p}\right)} - \frac{1}{\left(f_m^2 - f^2 + \frac{if_m f}{Q_m}\right)} \right) \end{aligned} \tag{4.10}$$

Here f_p, Q_p are the normal mode frequency and Q of the 'plus' mode (the higher frequency mode in the bar-transducer system) and f_m, Q_m refer to the 'minus' mode. These quantities can be extracted directly from the raw detector output as described in Sec. 5.2.4. The factor α has dimensions of inverse mass. In our mechanical model of Sec. 2 it would have a value $\frac{1}{2\sqrt{m_3 m_4}}$. In the end though, this value is folded into an overall scale factor which is determined experimentally as described in Sec. 5.2.4.

4.3 Displacement to current

What follows in the chain is the superconducting transducer. Motion of the secondary resonator pushes magnetic flux through a coil which induces a time varying current.

$$\begin{aligned} \tilde{I}(f) &= K \cdot \tilde{x}(f) \\ \tilde{I}(f) &= \left(\frac{I_{\rm DC}}{g}\right) \tilde{x}(f) \\ K &= \left(\frac{I_{\rm DC}}{g}\right) \end{aligned} \tag{4.11}$$

Here $I_{\rm DC}$ is a persistent current stored in the transducer, which is measured at the time that it is injected. The effective gap between the secondary resonator and the superconducting coil, g, is a parameter that can be determined by independent measurements of the transducer made prior to its installation on the bar (`https://sam.phys.lsu.edu/ALLEGRO/552`). However, the value of K can be folded into an overall scale factor which is determined experimentally as described in Sec. 5.2.4.

4.4 Current to voltage

For the DC SQUID, input current passes through a coil and produces flux – which then gives an output voltage. We fold the gain of a room temperature amplifier into the transfer function. The result is a simple constant which has the dimensions of an impedance.

$$\tilde{v}_S(f) = Z \cdot \tilde{I}(f) \tag{4.12}$$

The parameter Z can be determined from properties of the SQUID, and it will scale with the SQUID amplifier setting. But again, the value of Z can be folded into an overall scale factor which is determined experimentally as described in Sec. 5.2.4.

4.5 Demodulation

The lock-in amplifier and anti-aliasing low-pass filters are taken together. This step demodulates the signal and also introduces a time delay and phase shift to the signal. We treat

the output as a complex number, defining $z_v(t) = v_2(t) - iv_1(t)$. Care must be taken to preserve the phase of the signal at this point. The output depends on the relative phase of the signal and the lock-in reference oscillator, so the phase of this reference with respect to the time base of the measurement must be known. This is set so that it has zero phase (positive-going zero crossing) at each GPS 1 pps pulse. The reference frequency f_r is set to an integer number of Hertz so that this is possible. The convention chosen above preserves the phase – Fourier components are shifted in frequency by the demodulation but their phase is unchanged. We have

$$\begin{aligned} \tilde{z}_v(f - f_r) &= A(f, f_r)\tilde{v}_S(f) \\ \tilde{z}_v(f - f_r) &= a_L e^{i(t_d 2\pi(f-f_r))} e^{-i\phi_L} \tilde{v}_S(f) \end{aligned} \tag{4.13}$$

The overall amplifier gain a_L is determined, and will vary with the amplifier setting. Again this is folded into an overall scale factor which is determined experimentally as described in Sec. 5.2.4. The reference frequency f_r is set by the reference oscillator. The phase shift ϕ_L and a time delay t_d are parameters determined by measurements discussed in Sec. 5.1

Lastly the two analog voltages (real and imaginary part of z_v) are digitized. There is a scale associated with the discretization D – units of counts per volt. The discrete time series is represented here by z_c

$$\tilde{z}_c(f - f_r) = D\tilde{z}_v(f - f_r) \tag{4.14}$$

The constant D is a known parameter of the data acquisition system.

4.6 Overall response function

The overall response of the detector is thus represented as a product of these six transfer functions and we recover a frequency domain representation of the equivalent strain applied to the bar.

$$\tilde{h}(f - f_r) = \frac{\tilde{z}_c(f - f_r)}{JGKZAD} \tag{4.15}$$

We do not undo the demodulation, so the final output of the back-transformation are the demodulated frequency components of the strain. An inverse Fourier transform of $\tilde{h}$ produces a complex heterodyned time series $h^H(t)$ which has information in a 250Hz band centered at the reference frequency (near 900 Hz) at the 250 samples/s rate.

4.7 Calibrator actuation constant

Finally, we look at the path of the calibration signal, which shares the same path with a gravitational wave signal starting with the mechanical response of the bar to a force. Since we cannot produce a calibration gravitational wave in our laboratory that is strong enough to detect, we apply a force to the end of the bar. We already know how to relate a strain

to such a force through Eq. 4.9 – though it is apparent that this relationship cannot be verified with a direct measurement. The force is applied via the calibrator discussed in Sec. 2. Given the transfer function of the calibrator H we can apply a known force to the end of the bar and directly measure the response $GKZAD$. The calibrator was designed to have a very simple geometry and mechanics. In principle, knowledge of the capacitance and the area of the plates would allow us to calculate the actuation (applied voltage to force) constant (see Sec. 5.2.1). It was discovered that this simple model of the force generator is inadequate. Reciprocal measurements, where excitation and response were both done through the calibrator, allow a determination of H. Details are discussed in Sections 5.2.2 and 5.2.3.

5 Measurements

5.1 Timing

The goal is to produce a strain time series that has the correct phase relation with respect to a time scale that will allow comparison to signals from other detectors. There are some special considerations for the timing that arise for data that is demodulated through mixing with a reference signal in a lock-in amplifier prior to sampling. In the signal chain prior to the mixing, the information is passed at a higher frequency – in this case a band around ~ 900 Hz. A time delay in the signal chain before the mixing will have a different effect than one that occurs after the heterodyning. For example, a 30 μs delay before the heterodyne shifts the phase of 904 Hz by 9.8°. The range of the phase shift over a 100 Hz band centered at 904 Hz is only about a degree. A delay of this nature thus could be approximated by a constant phase shift over the relevant band.

After the mixing the information is passed in the relevant (~ 100 Hz) band around DC. In this case a 500 μs delay gives a 9° phase shift at 50 Hz – and a nearly 20° span over the relevant 100 Hz signal band (-50 to $+50$ Hz). This sort of delay is not well approximated as a constant phase shift and must be treated as a frequency dependent phase shift.

For the data acquisition system it is the reference oscillator that provides the timing before the mixing. This is why we synchronize the oscillator to GPS. It is the sample time of the DAQ that provides the timing for the low frequency demodulated signal after the heterodyne.

5.1.1 reference oscillator

The phase of the 904 Hz reference oscillator, an SRS model DS345 function generator, was reset three times just prior to and during S4. (For a list of resets see `https://sam.phys.lsu.edu/ALLEGRO/1097`, `https://sam.phys.lsu.edu/ALLEGRO/1223`, `https://sam.phys.lsu.edu/ALLEGRO/1238`)

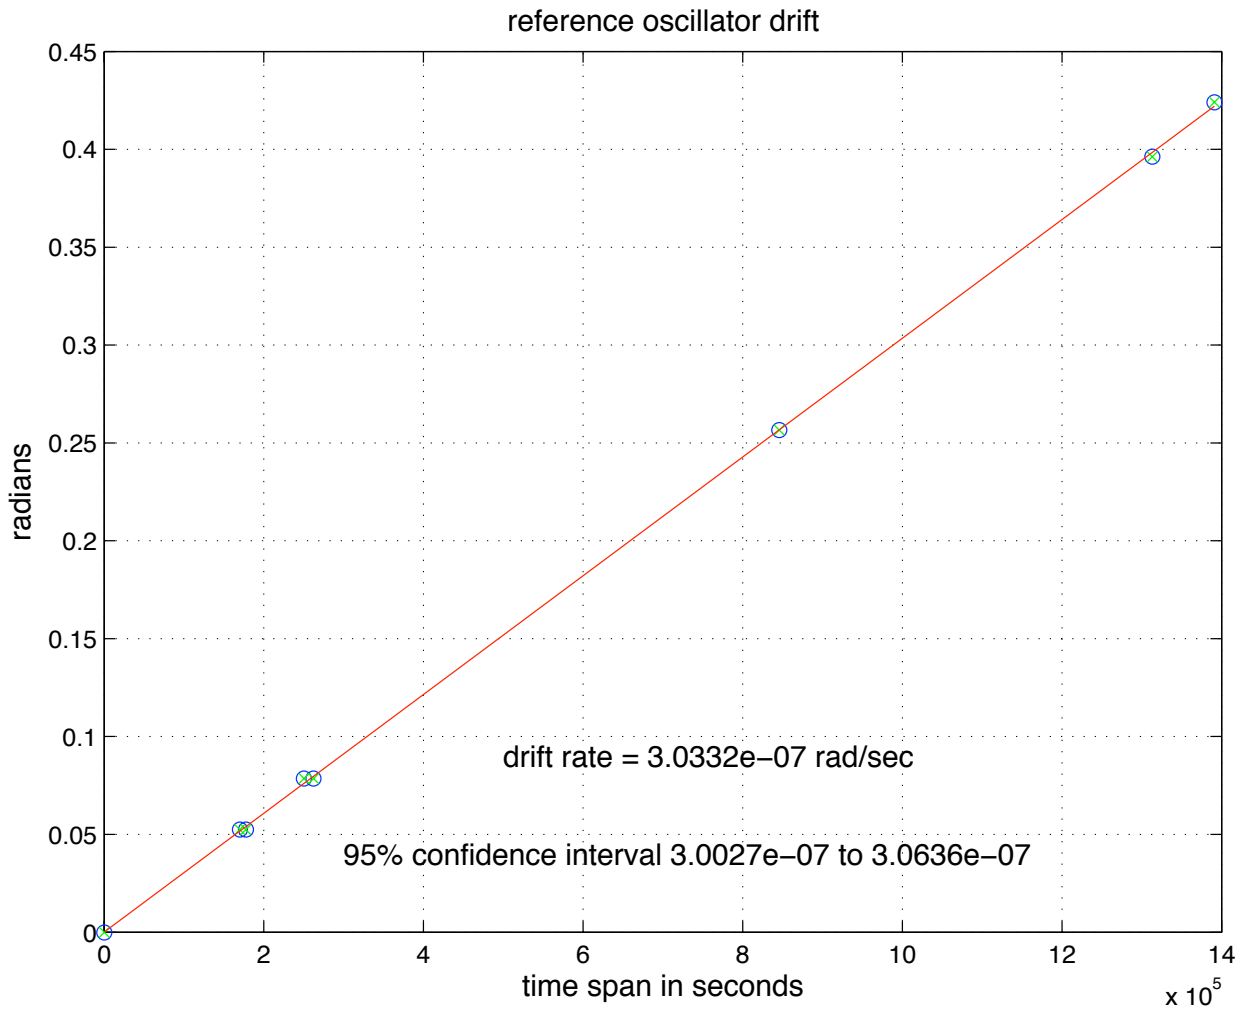

Figure 3: Plot of measurements of the reference oscillator drift during and around S4. The red line is a fit to a linear drift rate, which shows that the osciallator, which is a frequency synthesizer, is offset from the desired frequency by 48 nHz.

Each reset moved the sinusoid zero crossing to occur at the start of the GPS second. Looking at the amount of drift for these and several other resets done before and after S4, we see that they indicate a steady phase drift of magnitude 3.03×10^{-7} rad/s. Figure 3 shows the data points with the fitted slope. This drift rate, along with the time since the last phase reset of the oscillator, was used to determine the reference oscillator phase (ϕ_L in Eq. 4.13) for all times during the data run.

5.1.2 DAQ timing

When we attempted to understand the stochastic hardware injections in detail, it was discovered that the data acquisition system (DAQ) had a significant absolute offset between the nominal sample time and the actual sample time, an offset that varied from one injection to the next. Due to some fortunate foresight (thanks Bill Hamilton!), a DAQ channel had

been dedicated to recording a GPS 1 pulse-per-second (1pps) signal. Using this pulse, it was found that each launch of the DAQ (itself done with a GPS 1pps pulse) had an uncertainty of up to one sample time (4 ms in this case) (see `https://sam.phys.lsu.edu/ALLEGRO/1436`). Figure 4 shows the recorded 1pps signal for three different 'timing epochs' between which the DAQ was restarted. It appears as if the initial "trigger" input to our DAQ never actually started the sampling sequence, but instead acted as a "gate" to allow the next pulse in an unsynchronized pulse sequence to become the first in that epoch.

During a continuous run, or timing epoch, the stability of the DAQ timing was much better (see `https://sam.phys.lsu.edu/Data_Analysis/307`). The histogram plot reproduced here as Fig. 5 shows the spread of times extracted from the 1 pps. The stability is seen to be within a few μsec.

Therefore each restart of the DAQ produces a new timing epoch, where the timing stability is satisfactory, but the overall absolute timing needs to be determined. The absolute timing of a particular epoch was determined in two sets of measurements. First, by injecting a sine wave, whose zero crossing had been synchronized to the GPS 1pps, directly into a DAQ channel (`https://sam.phys.lsu.edu/Data_Analysis/276`). Second, the hardware injection system was used to inject a known broad-band signal at a known time (see `https://sam.phys.lsu.edu/Data_Analysis/282`). Using these epochs of known absolute timing we can go back and determine the absolute timing of any epoch using a relative timing offset. That relative timing offset is obtained via the recorded GPS 1pps channel and an interpolation method (see `https://sam.phys.lsu.edu/Data_Analysis/275`).

All of this leads to a correction offset of 1.24 ms for the timing during the S4 run (see `https://sam.phys.lsu.edu/Data_Analysis/281`). The sign is such that we should subtract 1.24 ms from the nominal time labels. In practice this timing offset is applied via a phase factor in the frequency domain by the calibration routines (see Sec. 7).

5.1.3 filter delays

Once the timing of the DAQ is known we can then turn to the time delays associated with any filters in the signal path.

The input filters of the lock-in produce a nearly frequency independent phase shift over our frequency band of interest. The output lock-in filters and the anti-aliasing filter that follows introduce a significant time delay. These values are determined in measurements of the isolated electronics (see `https://sam.phys.lsu.edu/Data_Analysis/291`) and were determined to be a 10.9 ms delay and a 15° phase shift. In fact, we see from these data that a constant phase plus a term linear in frequency (the time delay) do no account for the entire effect on the phase due to these filters. Higher order terms (quadratic in frequency and so on) can be fit as well. Although this is a small effect, these terms were included (see `https://sam.phys.lsu.edu/Data_Analysis/319`). As with the DAQ correction, these factors are applied via a phase factor in the frequency domain by the calibration routines

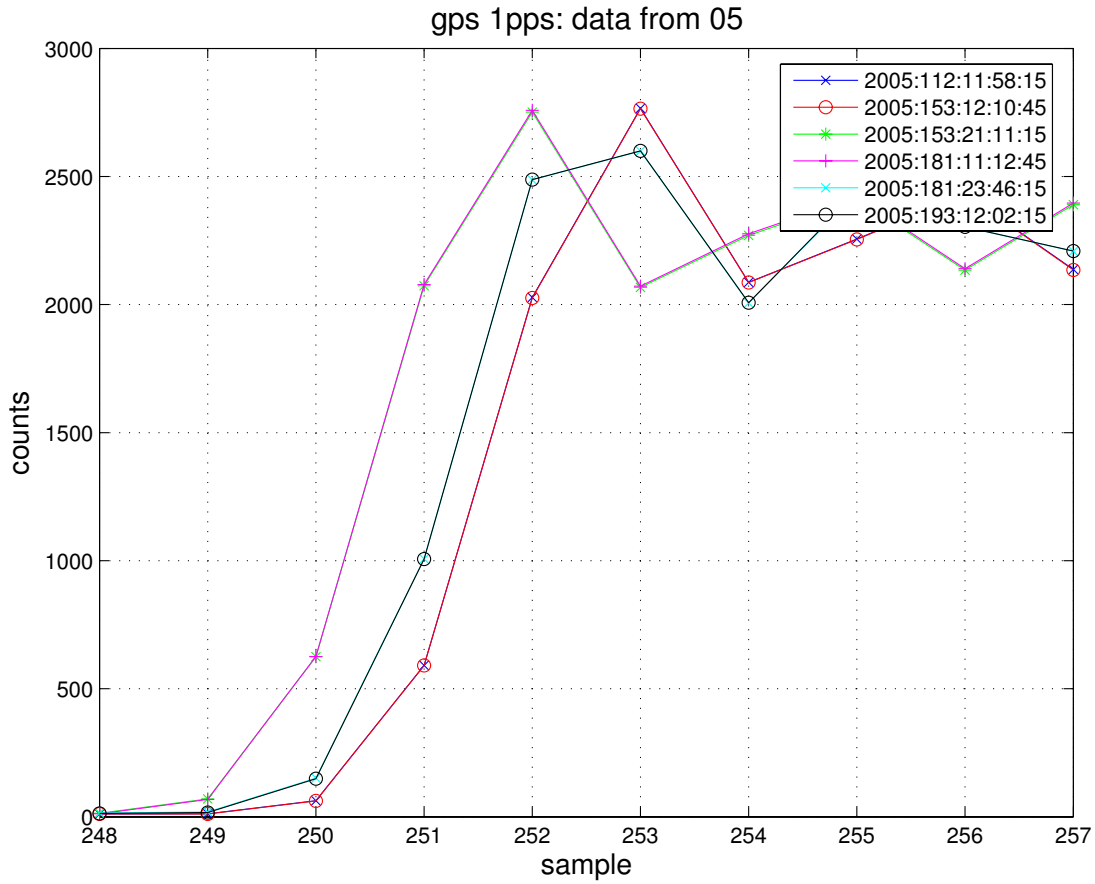

Figure 4: The figure shows fine detail at the leading edge of 6 recorded timing pulses, or "GPS 1pps" pulses – a pair for each of three different epochs in 2005. It can be seen the variation is miniscule during an epoch. There was a restart on day 153 at 19:37:15 and a restart on day 181 at 20:39:45. Each restart of the DAQ significantly shifts the recorded 1pps relative to the data point times. The horizontal axis shows a the sample number. The spacing between data points is 4 ms. (This plot is also available from the elog at `https://sam.phys.lsu.edu/ALLEGRO/1436`)

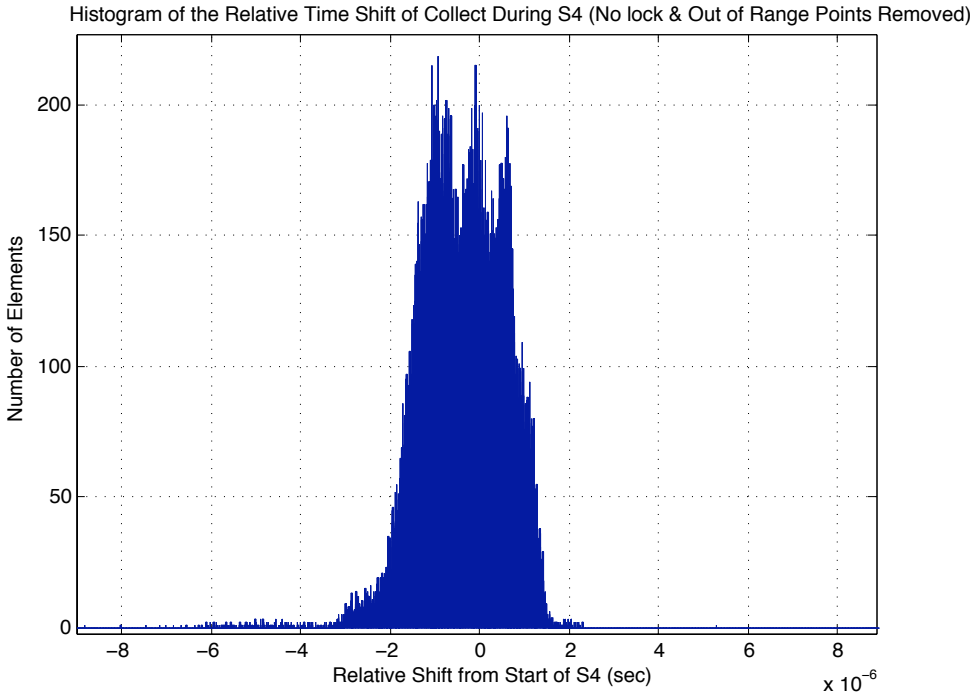

Figure 5: Histogram of the timeshift between the exact sample time and the leading edge of the GPS1pps pulse during one epoch of S4, showing that the time distribution function was about 4 microseconds wide.

(see Sec. 7).

5.1.4 Further checks

Tests were performed to check the timing accuracy of the hardware injection system. An input file was generated which produced a 1pps output that could be directly compared to the GPS clocks. The agreement between them was observed to be better than a microsecond (see `https://sam.phys.lsu.edu/ALLEGRO/1453`).

Additionally, as a sanity check, several sets of tests were done to compare the two independent GPS clocks (an HP 58503A and a TruTime) used in the system. They agreed to better than a microsecond. See `https://sam.phys.lsu.edu/ALLEGRO/1442`, `https://sam.phys.lsu.edu/ALLEGRO/1443` and `https://sam.phys.lsu.edu/ALLEGRO/1447`)

5.2 Magnitude

The above discussion on timing relates to producing $h(t)$ with the proper phase. We know turn to the task of producing it with the correct magnitude.

5.2.1 The force generator geometry

The force generator, or force actuator, or calibrator, or calibration capacitor, is shown in Fig. 6. The working parts of the capacitor are accurately to scale, the rest is a reconstruction, and is only approximate.

The capacitor plates themselves are both annular and matched in dimensions, with a nominal outer diameter of $d_o = 5.12$ inches and a nominal inner diameter of $d_i = 2.755$ inches, and an design gap g of 0.020 inches (see `https://sam.phys.lsu.edu/ALLEGRO/591`). Both are electrically isolated from ground and the bar by sapphire washers or epoxy, and each is connected to the outside by the center-conductor of a separate coaxial cable. Therefore we can make a special type of electrical measurement, a "three-terminal capacitance measurement using a ratio-transformer bridge". Such a system can measure the direct capacitance C_0 between the two plates, and leave out all of the capacitance to ground that inevitably comes with the cabling.

After assembly, this capacitance was measured to be 163.06 pF at room temperature, `https://sam.phys.lsu.edu/ALLEGRO/688`, and after the S4 run, while still cold, was found to be 164.04 pF, `https://sam.phys.lsu.edu/ALLEGRO/1310`. Since stray capacitance is not included in the measurement, to a good approximation this capacitance will be

$$C_0 = \epsilon_0 \frac{A}{g} \tag{5.1}$$

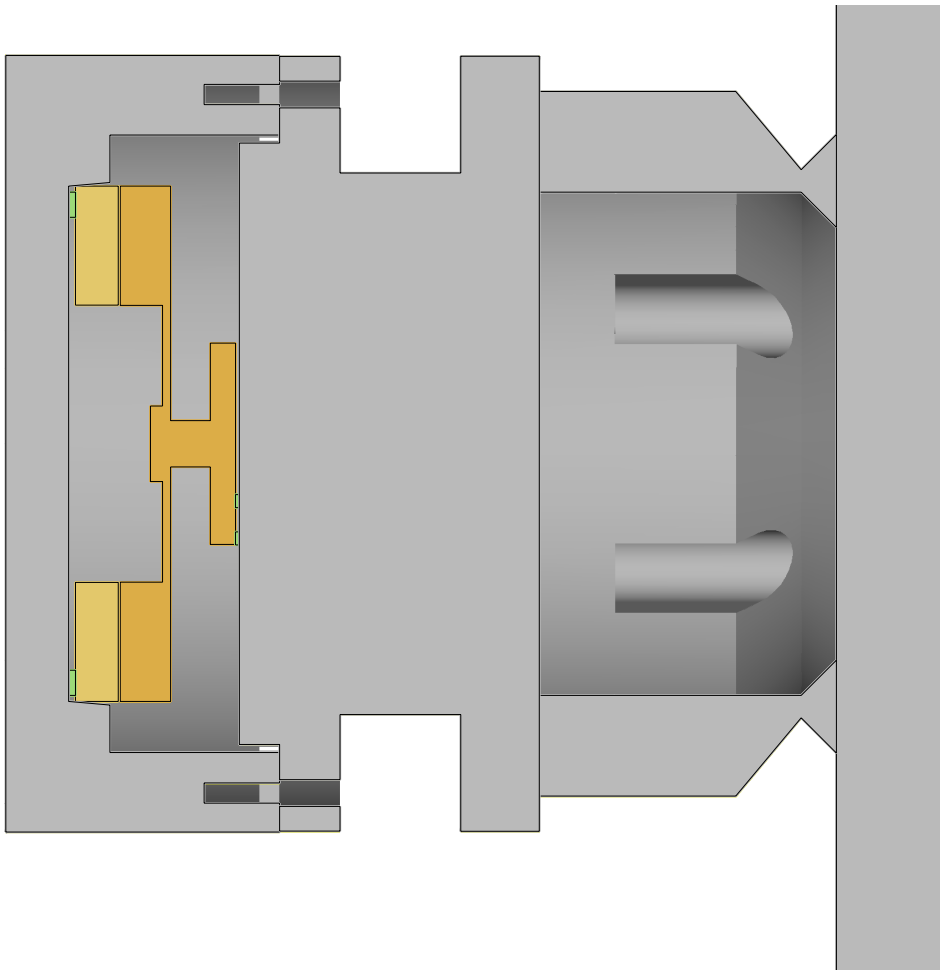

Figure 6: The calibrator, shown in cross-section, drawn to scale. The axis of the bar is horizontal and coaxial with the calibrator, and a small portion of the bar face is shown on the right. The lighter yellow capacitor plate is called the "outer plate", and the darker yellow one is called the "inner plate"; both plates are electrically isolated by insulators, shown in green. The gray area is the electrically grounded enclosure. The outer plate is mechanically connected to the "end cap" by a wide layer of epoxy (not shown), and so is rigidly attached to the end-cap of the enclosure. The inner plate is attached to its mounting pedestal only by thin flexures, so that its mechanical coupling to the bar is very weak. Pictures of it can be seen at `https://sam.phys.lsu.edu/ALLEGRO/593`.

where A is the annular area and g it the capacitor gap. Then we can infer the capacitor gap g (when cold) to be

$$g = \epsilon_0 \frac{\frac{\pi}{4}(d_o^2 - d_i^2)}{C_0} = 5.043 \times 10^{-4} \text{ m} = 19.9 \times 10^{-3} \text{ in} \tag{5.2}$$

where we have assumed a 0.5% linear contraction of the plate diameters due to the temperature change.

From this geometry we can directly calculate the AC electrostatic force $F_e(t)$ applied to the plates if they are have a combined DC and AC voltage difference between them of $V_{DC} - v_e(t)$. *Note the sign convention.* The instantaneous electric force F_e is

$$F_e = \frac{1}{2}\epsilon_0 \int E^2 dA \tag{5.3}$$

where E is the electric field, so, assuming a constant gap,

$$F_e = \frac{1}{2}\epsilon_0 \frac{V_{DC}^2 - 2V_{DC}v_e(t) + v_e^2(t)}{g^2} A \tag{5.4}$$

and if $v_e(t)$ is band-limited, and/or much smaller than $V_D C$, we have accurately

$$F_e(t) = -\left[\frac{\epsilon_0 A}{g^2} V_{DC}\right] v_e(t) \equiv H v_e(t) \tag{5.5}$$

and we have apparently found our force generator constant H

$$H = \left(-3.253 \times 10^{-7} \frac{\text{N}}{\text{V}^2}\right) V_{DC} \tag{5.6}$$

The one serious uncertainty left is the degree of tilt, or systematic variation in the gap as a function of azimuth angle. To first order, such variations must average out, but we don't know how to quantify any second-order effects.

Bound charge effect After several measurements of the transfer function from $v_e(t)$ to output z, we found that they were not consistent between positive and negative values of $V_D C$. See, for example, `https://sam.phys.lsu.edu/ALLEGRO/1219`. This turned out to be explicable in terms of a bound (or fixed) charge on the capacitor plates, whose value can be parameterized with an offset voltage V_{off}, so now Eq. 5.6 becomes

$$H = \left(-3.253 \times 10^{-7} \frac{\text{N}}{\text{V}^2}\right) (V_{DC} - V_{off}) \tag{5.7}$$

where $V_{off} = 1\text{V}$. We made no great efforts to determine the time variation of V_{off}, because it was a such small correction for the measurements done at high values of V_{DC}.

5.2.2 The reciprocity measurements

Introduction After the S4 run, quite a bit of effort went into making the so-called "gamma measurements". The ALLEGRO elog shows 25 "gamma runs", each involving a number of separate determinations of the various gammas.

The goal was to confirm the absolute value of the actuator constant H from Eq. 5.7. This constant plays the same role as the "standard source" plays for a radio receiver. When we have a standard source, we don't need to understand every link in the chain of components that make up the completer receiver, instead we need only to measure the output for this "standard" input, and we have determined the absolute response function, or transfer function.

There was an unexplained discrepancy in the early results (the "first trials"). This discrepancy was that the H measured near 200 Hz was $\tilde{5}\%$ lower than the "geometric" value, but $\tilde{2}5\%$ higher near 900 Hz. A refined understanding of the mechanics was needed (see Sec 5.2.3). This note represents an explanation of this resolution of the earlier discrepancy. A second series of trials to measure the gammas was undertaken as a further check for systematic and statistical errors, but the earlier results were fully confirmed.

Gamma measurement procedure The concept of the gamma measurement is to make use of the "reciprocal" property of many "transducers", such as our force actuator. Reciprocity says that an actuator is also a position sensor, and that it has the "**same**" coefficient for both directions of operation, if those measurements obey certain conditions. Evidently our capacitive actuator/sensor/transducer falls in this category.

A schematic version of the experiment is shown in Fig. 7. The normal configuration of the transducer is altered, as shown, so that one plate can be quickly changed from a voltage generator that applies an input to the capacitor, to a current-sensing amplifier that measures its output. During the entire operation the other (main) transducer is unaltered, except that the SQUID is turned off.

The procedure, as suggested in Fig. 7, has two distinct parts: pump, then sense. First, the sine-wave generator, producing a $v_e(t) = V_p \sin(2\pi f_1 t)$ exactly at the mechanical resonance frequency f_1, is switched on for a short time T, so that the mechanical resonator is pumped up. Now it has a displacement $x = x_p \cos(2\pi f_1 t)$, where the peak amplitude x_p is far higher than its thermal value. Then the switch is thrown, and this vibration causes a current $I_e(t) = I_p \cos(2\pi f_1 t)$ which is sensed by an "ammeter", also known as a "current amplifier". The current amplifier, by definition, senses current, but also has a "low enough" impedance, compared to the output impedance of the capacitor, so that it will not reduce the current from its maximum value. Then the result of each experiment can be parameterized by a gamma, *for this mode*, defined to be

$$\gamma_1 \equiv \frac{I_p}{V_p T} \tag{5.8}$$

which should be independent of the excitation time T and level V_p.

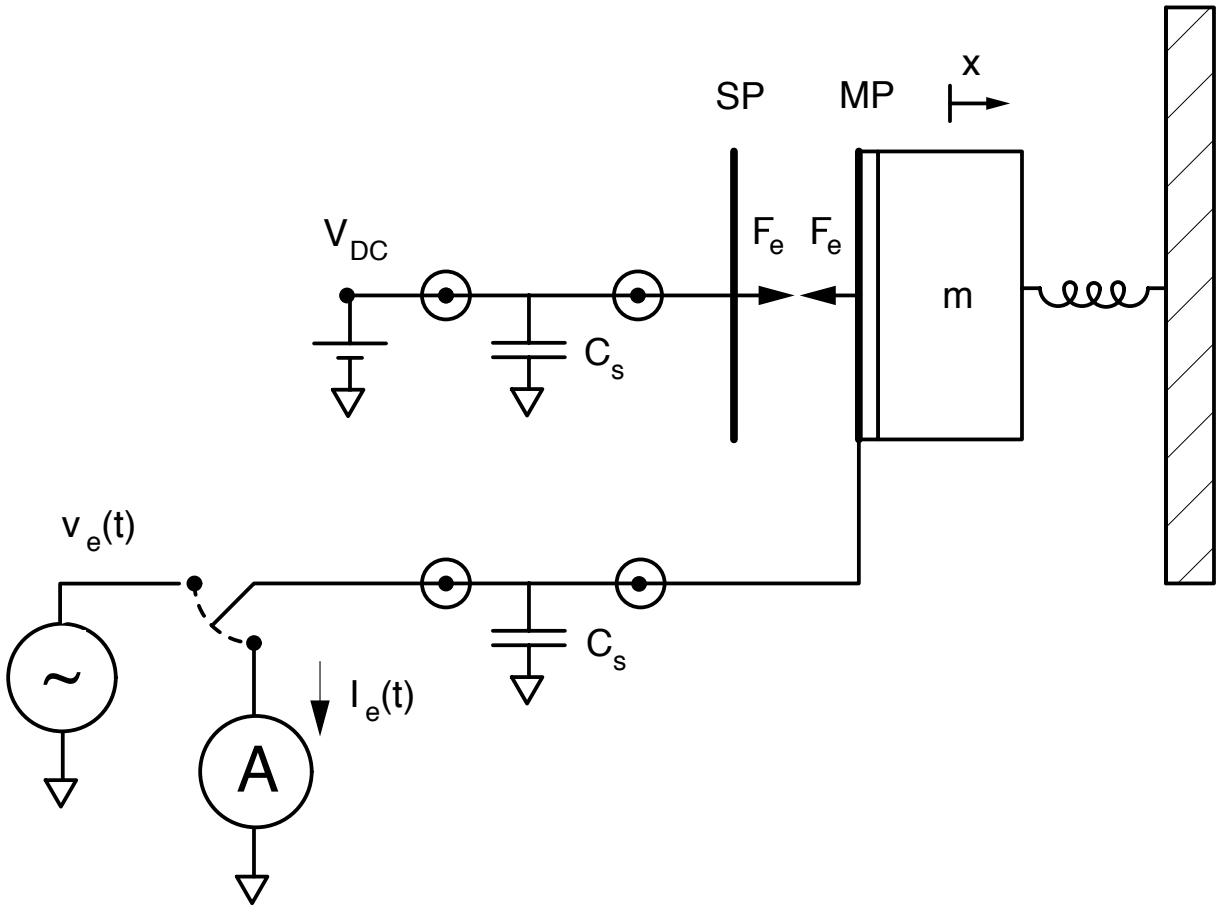

Figure 7: A simplified version of the "gamma procedure", which uses reciprocity to measure the actuator constant H. The moving capacitor plate MP is attached to the mechanical oscillator but is electrically insulated from it; the "stationary" plate SP (floating in inertial space) is biased by a battery to have voltage V_{DC}. The switch is first set to allow excitation exactly on resonance by the sine wave generator voltage $v_e(t)$ for a time T, and then is switched to the "ammeter" A and the amplitude of the pumped-up current $I_e(t)$ is measured. Then gamma is defined as the ratio of "output" to "input".

A simple relation between γ and H From the model in Fig. 7, we can then calculate how gamma is related to the actuation constant. The exciting force is defined, in terms of the actuation constant H, to be

$$F_e(t) = F_p \sin(2\pi f_1 t) = H V_p \sin(\omega_1 t) \tag{5.9}$$

The mechanical response function, in the frequency domain, is

$$\tilde{x}(\omega) = \frac{\tilde{F}_e(\omega)/m_1}{-\omega^2 + \omega_1^2} \tag{5.10}$$

where $\omega = 2\pi f$ is the fourier angular frequency and m_1 is the "effective mass" of the mode. This can be transformed to find the exact Green function in the time domain, which leads to a time response (neglecting the phase)

$$x(t) = \frac{F_p}{m}\frac{T}{2\omega_1}\cos(\omega_1 t), \quad t > T \tag{5.11}$$

or, in other words, the peak displacement caused by the pumping is

$$x_p = \frac{HV_pT}{2\omega_1 m_1} \tag{5.12}$$

Once the switch is moved over, we calculate the current from

$$I = \frac{dq}{dt} = \frac{dC}{dt}V_{DC} + \frac{dV}{dt}C_0 \tag{5.13}$$

where $V(t)$ is the AC voltage across the capacitor, and the second term on the right is the usual current voltage relationship for a capacitor. The first term can be thought of as a current generator I_e caused by changes in the capacitor gap $g + x$, i.e. by the motion of the mass,

$$I_e = V_{DC}\frac{dC}{dt} = V_{DC}\left(\frac{dC}{dx}\right)\left(\frac{dx}{dt}\right) \cong V_{DC}\left(\epsilon_0\frac{A}{g^2}\right)(\omega_1 x_p \sin(\omega_1 t)) \tag{5.14}$$

or the peak current, after pumping, is

$$I_p = \left[\epsilon_0\frac{A}{g^2}V_{DC}\right]\omega_1\left(\frac{HV_pT}{2\omega_1 m_1}\right) \tag{5.15}$$

so immediately we have

$$\gamma_1 = \frac{H^2}{2m_1} \tag{5.16}$$

the main result. This result has a simple interpretation: a factor of H comes from actuation, a factor of $1/(2m_1)$ comes from mechanics, and one factor of H comes from sensing. Note that the result depends upon the "mode mass" m_1.

Extra precautions As noted above, a second set of gamma measurements was made in the summer of 2006. Their primary purpose was to determine if there were flaws in the measurement procedure, and as a refinement and check on the earlier measurements.

The main change in configuration was a change in the plates where the DC and AC voltages were applied. In the first set of trials, we had often applied *both* voltages to the same plate, with the other plate grounded. The electric force between the plates (proportional to $(V_{DC} + v_{AC})^2$ was exactly as desired, but there must also have been electrostatic fields and forces with the same exact time dependence being applied to the 'back-side' of the biased plate, and to the enclosure wall behind it. These extra forces could conceivably have contributed, so we eliminated them in the second trials by adopting the voltage bias configuration shown in Fig. 7. The DC voltage now was applied to only one plate, the AC voltage only to the other. Therefore there was no region where extraneous forces at the excitation frequency

could exist. (See the circuit diagram attached to the gamma18 entry, for an example of separation of AC and DC voltages. https://sam.phys.lsu.edu/ALLEGRO/1510)

The measurement of duration T for the excitation phase was improved by using a "burst generator" for $v_e(t)$. It produced a burst with exactly the specified number of cycles of a sine wave at the desired frequency, thereby removing any errors associated with manual operation of a stopwatch for measuring the duration.

The current amplifier, a SR830 lockin with two current inputs, was tested for accuracy at the measurement frequency with a calibrated current source with an output impedance of 10^7 ohms. https://sam.phys.lsu.edu/ALLEGRO/1313. The input labeled $I(10^8)$ was not accurate, but the one labeled $I(10^6)$was. Since we used this setting for all measurements, we were reassured that all measurements of peak 'return' current I_p were accurate.

We found that the minus mode frequency was significantly "pulled" by the SQUID feedback ($\sim$ 80 mHz). This was learned in the middle of the first gamma series (about gamma5). In all the later measurements we took full precautions to determine the exact frequency for every mode with the SQUID off, so that we could be certain the pumping frequency was close enough to maintain good coherence throughout the pumping time T.

Finally, the software used to measure the time and phase dependence of the current I_p was modified to remove any bias due to DC offsets. It was also modified to determine the "early decay" correction, that is the effects of frictional loss during the excitation phase. These changes were only incorporated in the second trials analysis.

These later measurements were NOT substantially different from the earlier ones, and confirmed their results, within reasonable errors. As a result, no electronic or procedural error was found. The earlier measurement results are summarized here https://sam.phys.lsu.edu/Data_Analysis/212, and https://sam.phys.lsu.edu/Data_Analysis/294. Because the later measurements have better accuracy, and less possibility of systematic error, we report the details of those only. For a detailed summary of exactly where the data used for the second series came from, and how the final numbers were arrived at, see https://sam.phys.lsu.edu/ALLEGRO/1729.

Results for the "calibrator mode" Looking back at the cross-section drawing of the calibrator, Fig. 6, the inner plate has a low-frequency mode (of translation parallel to the axis) that is a good test of the entire gamma procedure. Because γ (see eq. 5.16) varies as the inverse of the "mode mass" m_1, this mode has a very large gamma, so is very easy to excite. It also turns out to have a very long decay time, proving that it has unusually low vibrational friction in its connection to the bar.

This inner plate resembles a wheel mounted by thin spokes to a pedestal; the pedestal is attached to the bar via electrical insulators. The plate's physical mass is estimated to be 318 grams, from measurements of its' dimensions; the "effective mass" for this translational mode will be essentially the same as the physical mass. The frequency of the translational

mode mode of vibration was designed to be 200 Hz, and was measured to be 201.5 Hz when cold.

Only a few gamma measurements were made of this mode, because its very long decay times cause a long delay until the antenna can return to observing mode. The second trials, at $V_{DC} = \pm 40$ V, gave identical values (difference $< 0.5\%$) of 2.36×10^{-10}mhos/s (orders of magnitude higher than the bar mode gammas, as expected). Using Eq. 5.16 and the mass above, we find $H = 3.06 \times 10^{-7} N/V^2 V_{DC}$. This is 5.3% lower than the value expected from the 3-terminal capacitance measurement of the gap.

For many purposes, this 5% difference is perfectly good agreement between independent methods. On the other hand, our current understanding of measurement errors suggests we should have gotten better agreement, perhaps $\sim 1\%$ difference. Perhaps this is an indicator of the size of the systematic error in the geometric calculation due to the 'non-flat' plates.

Results for the main modes Many more measurements were made of the gamma's for the "plus" and "minus" modes of the bar. These results are shown in Figs. 8 and 9. These measurements tell a pretty consistent story.

The minus mode values have significant scatter, perhaps because they are rather small, about 10 times smaller than for the plus mode. This difference in absolute values indicates that we are fairly far away from optimal tuning. This value of detuning is consistent with the value inferred from the "tuning curve" (see Appendix A).

The plus mode results are more sparse, because they take longer to measure. One value was measured twice, and the reproducibility was remarkably good. (The blue circle and the red plus fall exactly on top of each other in the figure.)

Modeling shows that we can combine the gammas for the two modes into a single value that is supposed to be insensitive to the tuning. In other words, we are supposed to replace the gamma in our initial simple model with the sum of the two gammas, and then use the bar's effective mass for m_1

$$H^2 = 2(M/2)(\gamma_+ + \gamma_-) \tag{5.17}$$

where $M/2$ is the effective mass of the bar, and M is the physical mass. As an aside we note that this result is in complete agreement with the 'mixing angle' formalism that we have used in the past. The angle describes the degree of tuning of the modes and gives the following for H:

$$\theta = \tan^{-1}\left(\sqrt{\frac{\gamma_-}{\gamma_+}}\right) \tag{5.18}$$

$$H = \frac{\sqrt{\gamma_- M}}{\sin\theta} \tag{5.19}$$

and

$$H = \frac{\sqrt{\gamma_+ M}}{\cos\theta} \tag{5.20}$$

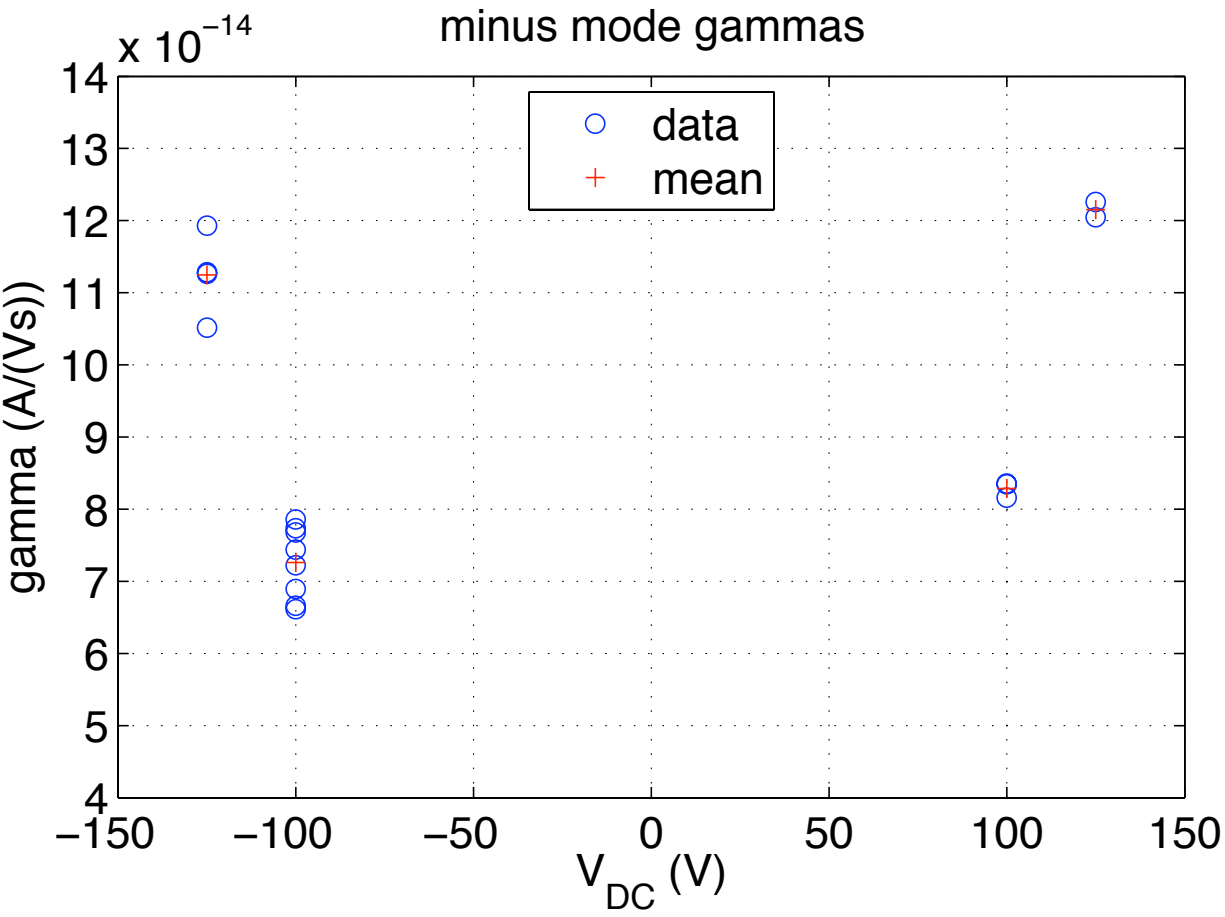

Figure 8: Results of measurements of the gamma ratio for the minus mode.

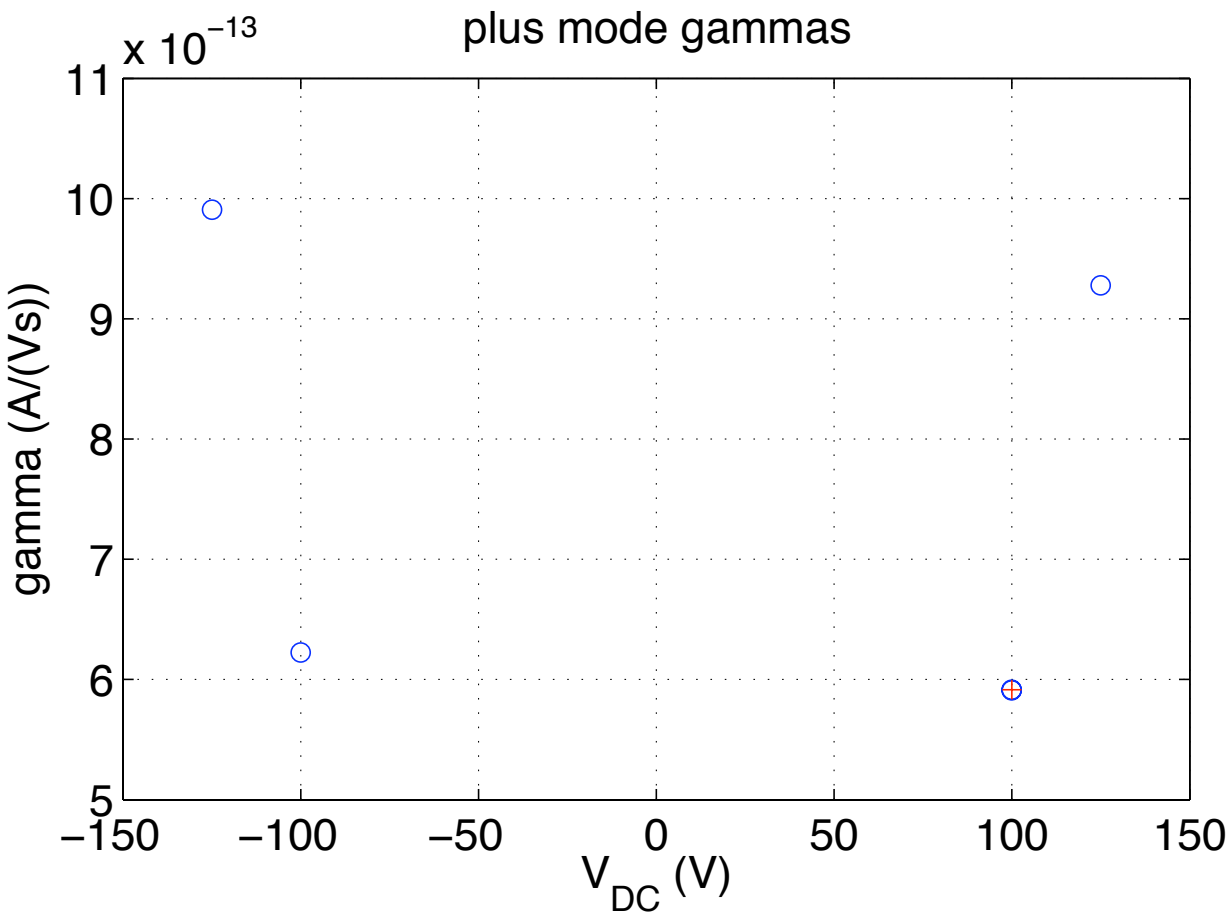

Figure 9: Results of measurements of the gamma ratio for the plus mode.

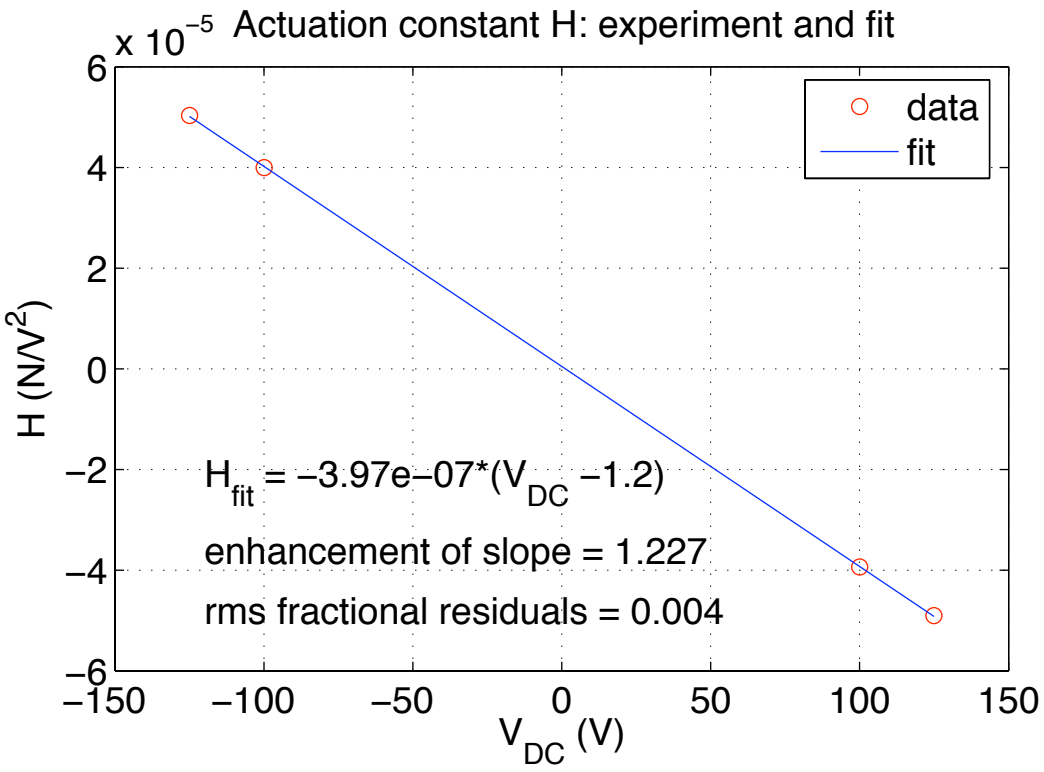

Figure 10: Results of measurements of the total gamma. Also shown is the best fit linear function. The 'enhancement' factor is the ratio of experimental slope to the slope calculated from the 'geometric' values of Eq. 5.7

Continuing with the measurements, each value of voltage provides an independent measure of H. The four values obtained are plotted in Fig. 10, and then fit a linear function. The results of the fit are also shown on that plot. (The data actually determine only H^2, so the overall minus sign for H is supposition.) The remarkable thing about this result is the enhancement (or increase) in the actuation constant. The enhancement here is 23% over the simple capacitor result, and 28% over the result obtained at 201 Hz.

Also remarkable is the finding of a voltage offset, as inferred by other measurements at lower voltages. It is comforting that they agree.

There is not enough data here to nail the uncertainty, but it is consistent with very small uncertainty, having a rms of the fractional residuals equal to 0.4%.

Again, comparing to the earlier determinations based on different data and an independent analysis (see `https://sam.phys.lsu.edu/Data_Analysis/212` and `https://sam.phys.lsu.edu/Data_Analysis/294`) we see the results are reasonably close. In both cases the force generator coefficient was determined to be ∼25% larger than expected based on geometry alone. This agreement between the two sets of measured results gives us considerable confidence in that result.

5.2.3 The interpretation of the reciprocal measurements

The full set of the gamma measurements all agree. So the question remains: how can the actuation constant be 23% high near 900 Hz, as compared to a "geometric" (expected) value?

So we have explored possible mechanical effects, to understand the results, and thus to determine which of the values to use near 900 Hz. Is it the one deduced from the measured geometry, or the enhanced value found by reciprocity measurements?

Normal mode mechanics The crucial mathematical tool for these investigations is the decomposition into normal modes. We restrict ourselves to motion in 1-dimension (≡x) for N masses, where each mass m_i is acted upon by external forces $f_i(t)$ in the x-direction. The masses are coupled by springs to each other and/or to mechanical ground. The math is described in some detail in Appendix B.

The final general result for the input/output matrix, Eq. B.13, can be simplified, by assuming only one mode a is excited, and only one force $\tilde{f}_i(\omega)$ is acting, then the input $\tilde{f}_i(\omega)$ produces an output $\tilde{x}_j(\omega)$

$$\tilde{x}_j(\omega) = \left[m_j^{-\frac{1}{2}} V_{ja} \frac{1}{(-\omega^2 + \omega_a^2)} V_{ia} m_i^{-\frac{1}{2}} \right] \tilde{f}_i(\omega) \tag{5.21}$$

Reading from right to left, this expression can readily be interpreted:

$$V_{ia} m_i^{-\frac{1}{2}} \tag{5.22}$$

is the factor that converts force applied to i into the amplitude of mode a, and

$$m_j^{-\frac{1}{2}} V_{ja} \tag{5.23}$$

is the factor that converts mode amplitude of a to displacement of j.

A force 'enhancer' ? We can use the expression above to calculate the 'enhancement' (or attenuation) of a force, depending on where it is applied. Consider a simple harmonic oscillator of mass m_2 and spring k_2, which has a small mass m_1 attached to the its face with a *stiff* spring k_1, as shown in Fig. 11.

We consider only the lower normal mode, or bar mode, and label it a. We are interested in "case-2", where we apply a force to the face of the bar m_2, and "case-1", where we apply the force to the small coupling mass m_1, which is also attached to the face. For the same external force, the ratio of excitation (case-1/case-2) is

$$\eta = \frac{V_{1a} m_1^{-\frac{1}{2}}}{V_{2a} m_2^{-\frac{1}{2}}} \tag{5.24}$$

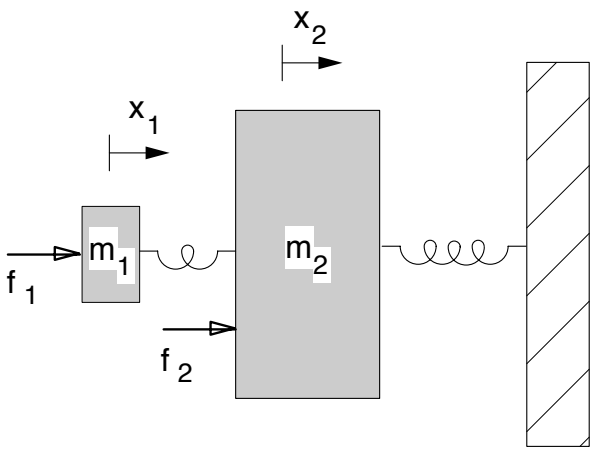

Figure 11: Model for possible force enhancer.

To make the example closer to ALLEGRO, pick $m_1 = 11$kg, $m_2 = 2296/2$, $f_2 = 900$Hz, and quantify the coupler stiffness with a frequency f_1, where $k_1 = m_1(2\pi f_1)^2$.

The resulting η is shown in Fig. 12. We do indeed find an enhancement in the force applied to the system. We find a coupler "frequency" of ∼2000 Hz would account for the experimental value of H derived from the gamma measurements. Physically what has happened is the the small mass m_1 is acting as an impedance matching device, and allowing somewhat more energy to flow to the bar. There is no conservation law for force the require it to be constant along a chain.

There are many modes observed in the vicinity of this plausibility argument. The mode observed here `https://sam.phys.lsu.edu/ALLEGRO/1403` that couples strongly to the calibrator is a likely candidate.

Sensing enhancement Examination of Eq.5.23 shows that there is an identical enhancement η in the sensing of the normal mode motion, and so the round trip measurement yields the following equation:

$$(\eta H)^2 = 2(M/2)(\gamma_+ + \gamma_-) \tag{5.25}$$

and so the values for H that we measured earlier from Eq. 5.17 are actually ηH. Luckily, this is just the quantity we really want.

5.2.4 Magnitude determination measurements

Now that we understand the calibrator (force generator), we can apply a known force to one end of the bar. The measured response will then allow us to determine the overall scale (or magnitude) of our data in terms of gravitational strain. Referring to Sec.4 we look at the

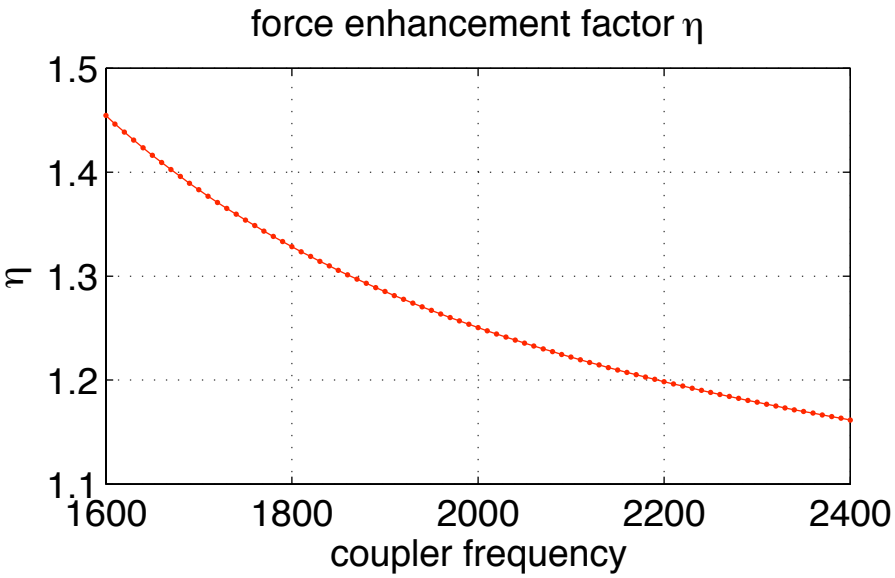

Figure 12: Result: a light coupling mass m_1 enhances the force applied to the bar.

full transfer function that takes us from calibrator input to recorded output.

$$\tilde{z}_c(f - f_r) = HGKZAD\tilde{v}_e(f) \tag{5.26}$$

Taking the H we determined above (really the enhanced ηH), in principle we just solve and get $GKZAD$. This then gives us the complete detector response. In practice we still treat some of the terms individually.

Mode frequencies Beginning with the frequency dependent terms we see in G, the mechanical response of the bar - secondary resonator system, we see that the response function requires the two mode frequencies and Q's. These are obtained directly from the raw data (see Fig. 13).

The response is extremely sensitive to the precise values of the mode frequencies, so it is necessary to track the frequencies as they slowly fluctuate due to small temperature changes. See also `https://sam.phys.lsu.edu/Data_Analysis/346` for an explanation of the mode tracking algorithm.

Figure 14 shows the two mode frequencies as they were tracked by the calibration code in 30 minute stretches over S4. The values of mode Q's are less critical – so average values over the entire S4 data run were used.

Frequency independent scale The factors that remain are all frequency independent constants and thus constitute an overall scale factor or gain. Some of the factors, such as the lock-in gain a_L and the digitization factor D can be determined independently quite easily. Others such as the mechanical gain α (see Eq. 4.10) and the effective gap g (see

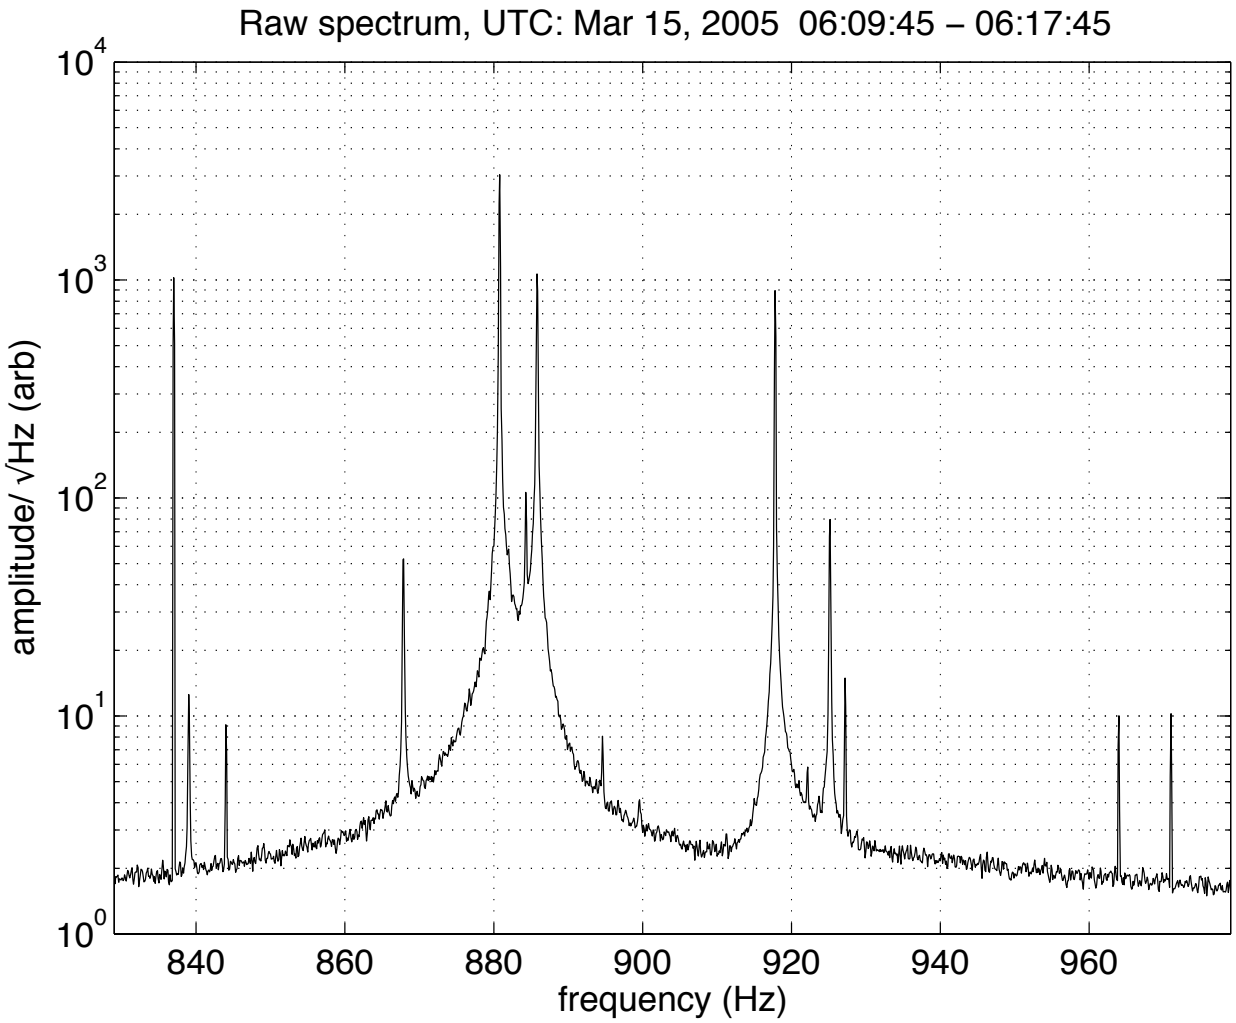

Figure 13: The figure shows the ALLEGRO raw amplitude spectral density (in counts$\sqrt{\text{Hz}}$) for 8 minutes of data taken during the LIGO S4 science run. The mechanical modes are the largest peaks at $f_m = 880.8$ Hz and $f_p = 917.8$ Hz. There is a calibration line (sinusoidal voltage applied to the calibrator) at 837 Hz. The other peaks are due to nuisance modes which add noise but do not couple significantly to a gravitational wave strain.

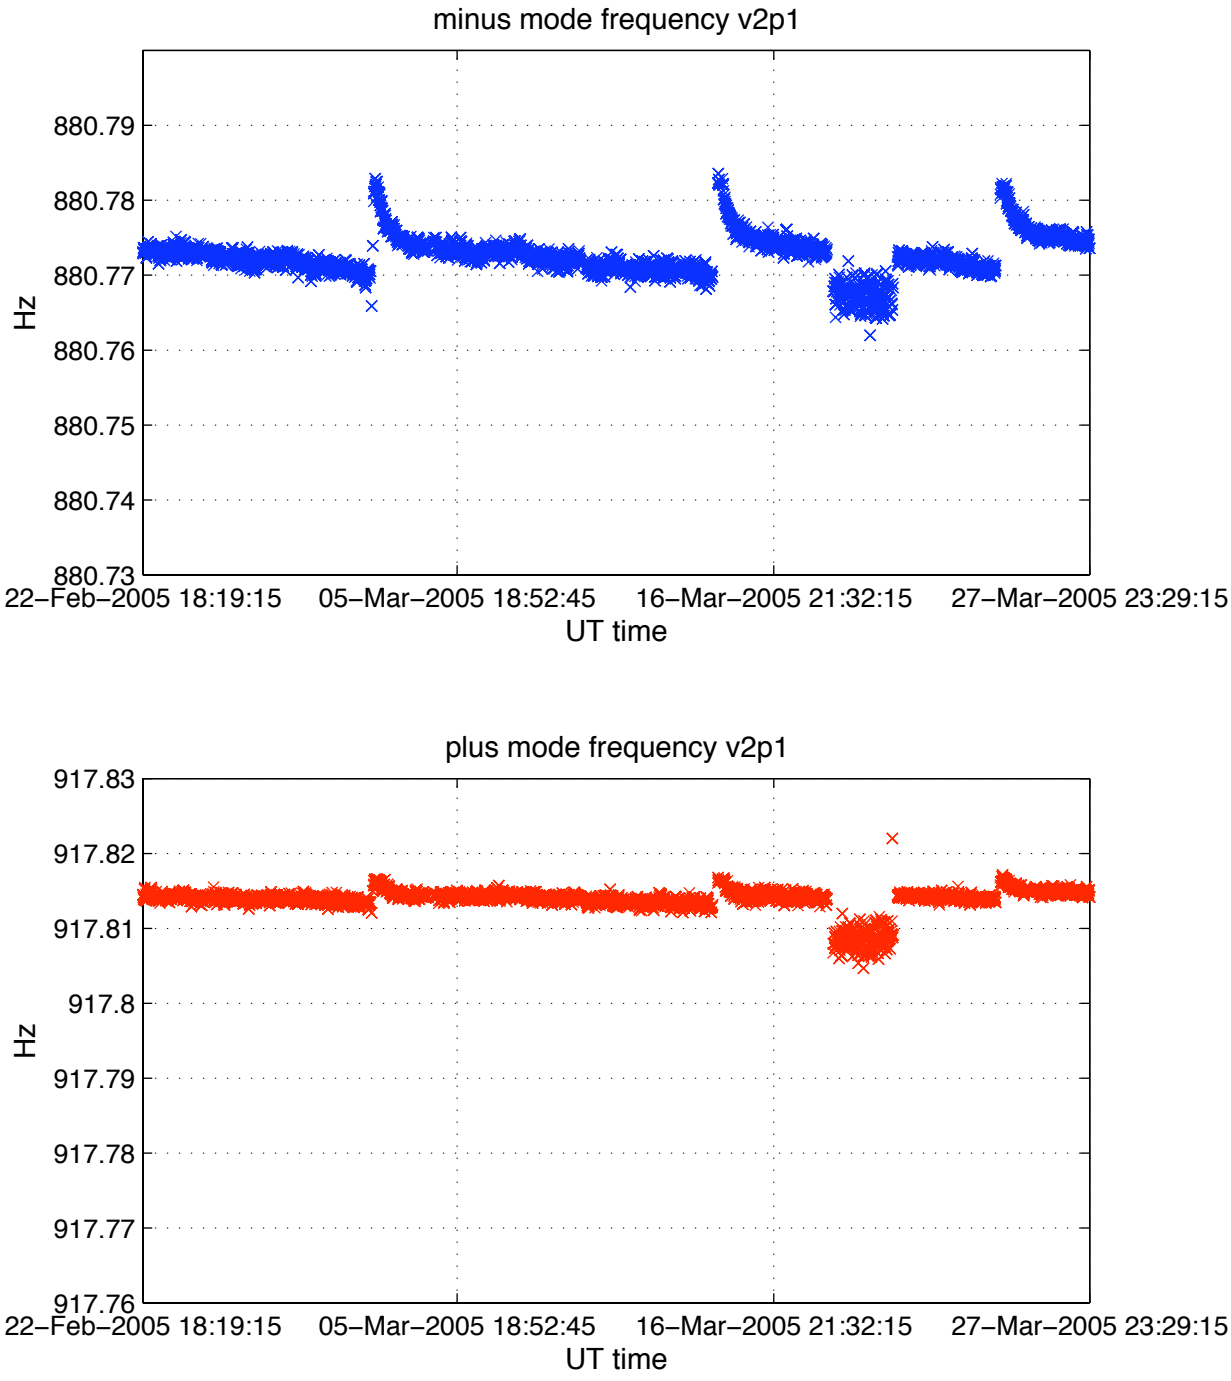

Figure 14: The figure shows the ALLEGRO normal mode frequencies tracked during S4. The three distinct jumps up correspond to LHe transfers. The jump down to a slightly noisy section, around March 20, occurred immediately following the rotation to the NULL orientation. The change is due to a higher level of feedback which was then corrected while still in the NULL orientation. (see `https://sam.phys.lsu.edu/ALLEGRO/1239`)

Eq. 4.11) of the transducer cannot be measured independently on a working detector. It will be preferable to measure this overall gain by introducing a calibration signal as alluded to above. This calibration procedure will serve at once to determine this overall scale and to verify the entire signal flow model. The procedure used for ALLEGRO is the following: a band-limited white noise signal is applied to the calibrator. First, the excitation signal is recorded by the usual DAQ through a separate lock-in. There are corrections for the small differences between this 'excitation lock-in' and regular 'signal lock-in' (see `https://sam.phys.lsu.edu/Data_Analysis/193`). With this correction we can then effectively divide through by $AD\tilde{v}_e(f)$

$$HGKZ = \frac{\tilde{z}_c(f - f_r)}{AD\tilde{v}_e(f)} \tag{5.27}$$

Note that for the division on the right hand side of Eq. 5.27 we must associate the frequencies in the excitation spectrum with the demodulated output by shifting by the reference frequency f_r. We divide by the frequency dependent part of G, denoted G_{freq} (see Eq. 4.10).

$$H\alpha KZ = \frac{\tilde{z}_c(f - f_r)}{G_{\text{freq}}AD\tilde{v}_e(f)} \tag{5.28}$$

Now, we divide by the calibrator transfer function – the force generator constant – H

$$\alpha KZ = \frac{\tilde{z}_c(f - f_r)}{HG_{\text{freq}}AD\tilde{v}_e(f)} \tag{5.29}$$

The right hand side is the number that the calibration procedure provides. It is convenient to provide nominal values for K and Z, but we do not have precise independent knowledge of these. Any error in these assumed value will be absorbed by the value for α We thus determine an overall gain for the gravitational wave detector. The results of this procedure are presented here: `https://sam.phys.lsu.edu/Data_Analysis/196`. We notice that even after the modeled frequency dependence is divided out there remains a small frequency dependent slope. One possible explanation is the effect of higher frequency mechanical modes (see `https://sam.phys.lsu.edu/Data_Analysis/211` and Appendix A). Given the uncertainty of this origin and complications in incorporating this into the model it was decided to leave the model unchanged. It was however determined that this unmodeled slope contributes to the overall uncertainty of the calibration.

Chronologically, these measurements were done before we had complete understanding of the calibrator. So the results need a correction factor of 1.24 – this is just the ratio of the calibrator actuation constant determined by the gamma's to that determined by the calibrator geometry (see Sec. 5.2.3). This corrected factor is implemented in the parameters file (S4params.m) as the factor **bar_trans_gain**. See Sec. 7 for more on the implementation.

Finally, the lock-in gain is not completely frequency independent, but was found to have a slight roll-off at higher frequencies. This is a small correction described here: `https://sam.phys.lsu.edu/Data_Analysis/296`.

5.3 The nuisance mode

Figure 13 shows an evident peak at 885.8 Hz. This very low Q, noisy mode was dubbed the 'nuisance mode'. The question arises as to whether this mode, which is highly excited by injections through the calibrator, should be considered as a 'detection' mode. There is good evidence that this mode is not strongly coupled to gravitational strain, but is strongly coupled to the calibration force generator. First, the tuning curve (Appendix A) which describes the mechanics of the readout transducer, is well-modeled without this extra mode. Second, the gamma for the nuisance mode is larger even than the plus mode, even though it is less well tuned to the fundamental bar frequency. This suggests a small effective mass for a mode as it is excited by the calibrator. Finally, the gamma measurements described in the previous sections were done at three different stored currents – three different tunings. The nuisance mode shifted much less in frequency than the plus and minus modes (see `https://sam.phys.lsu.edu/ALLEGRO/1508`). This again points to this not being a mode of the main readout transducer. We are able to notch out a band around this mode and ignore it in the calibration procedure. We are confident that not treating it as a detection mode is the proper thing to do. As seen in Fig. 17, in terms of strain, the peak for this mode thus becomes a noisy band.

5.4 Overall sign

It is crucial to get the overall sign of the response function correct, and there are multiple sign flips along the signal chain. The determination of the overall sign is discussed in detail here: `https://sam.phys.lsu.edu/Data_Analysis/293`

5.5 Hardware injections

See the following elog entries for descriptions of the analysis of hardware injections for a single detector `https://sam.phys.lsu.edu/Data_Analysis/320` and `https://sam.phys.lsu.edu/Data_Analysis/321` An important check is the ability to extract an injected signal with the appropriate phase – and understanding the injection system is essential for determining what 'appropriate' means. For ALLEGRO the dominant effect that needed to be accounted for was the 'sample and hold' of the D/A portion of the injection system. See `https://sam.phys.lsu.edu/Data_Analysis/290` and `https://sam.phys.lsu.edu/Data_Analysis/282`

The same type of effect is present for LIGO injections as well. It can be summarized by noting that the actuation function is the transfer function between an injection file and the mirror motion and so this must be accounted for in determining the actual injected strain signal. `http://ldas-sw.ligo.caltech.edu/ilog/pub/ilog.cgi?group=stochastic&date_to_view=12/06/2005&anchor_to_scroll_to=2005:12:30:14:19:55-mchugh`

5.6 Stability

Fig. 15 shows the stability of the calibration line in the version 2.1 calibrated S4 data. This is a constant voltage amplitude sinusoid applied to the calibrator continuously over the run. The stability of the output is better than a couple of percent. See the figure caption for details.

6 Uncertainties

6.1 Magnitude

The overall magnitude uncertainty will include contributions from the uncertainty in the known input and the uncertainty in the measured output. The input level amounts to knowing the force generator constant (H) of the calibrator, as well as knowing the DC bias and AC excitation voltages. We estimate its uncertainty at 2%. Output level uncertainty is given at 7% in `https://sam.phys.lsu.edu/Data_Analysis/196`. This is the statistical uncertainty coming from the spread in the measurements. There is also an unmodeled slope in the scale factor. We chose the value from a single point at the center of the sensitive band of the detector. This is at the 'plus' mode near 917.8 Hz. We estimate there could still be up to ±5% systematic uncertainty over the whole detection band arising from this unmodeled slope.

Combining 5%, 7% and 2% in quadrature we conservatively get a calibration magnitude uncertainty of ±10% The stability of the calibration as seen in `https://sam.phys.lsu.edu/Data_Analysis/309` is better than a couple of percent.

6.2 Phase

Again we have two distinct cases – pre and post heterodyning timing uncertainty. The first can be well approximated by a constant phase over our band. The second is gives a frequency dependent phase offset with a zero at the physical freq of 904Hz.

Pre heterodyne timing The reference oscillator is set to within 0.5 degrees. An estimate of the uncertain drift rate is $\pm 1 \times 10^{-8}$ rad/sec, which over the longest stretch without a reset gives only a 0.8° error in phase. Conservatively, we assign ±1° or ±3μs. In `https://sam.phys.lsu.edu/Data_Analysis/291`, an uncertainty of ±2° for the phase uncertainty. This comes primarily from lock-in input filter delays. To be very conservative the total phase uncertainty (pre heterodyne) is ±3°.

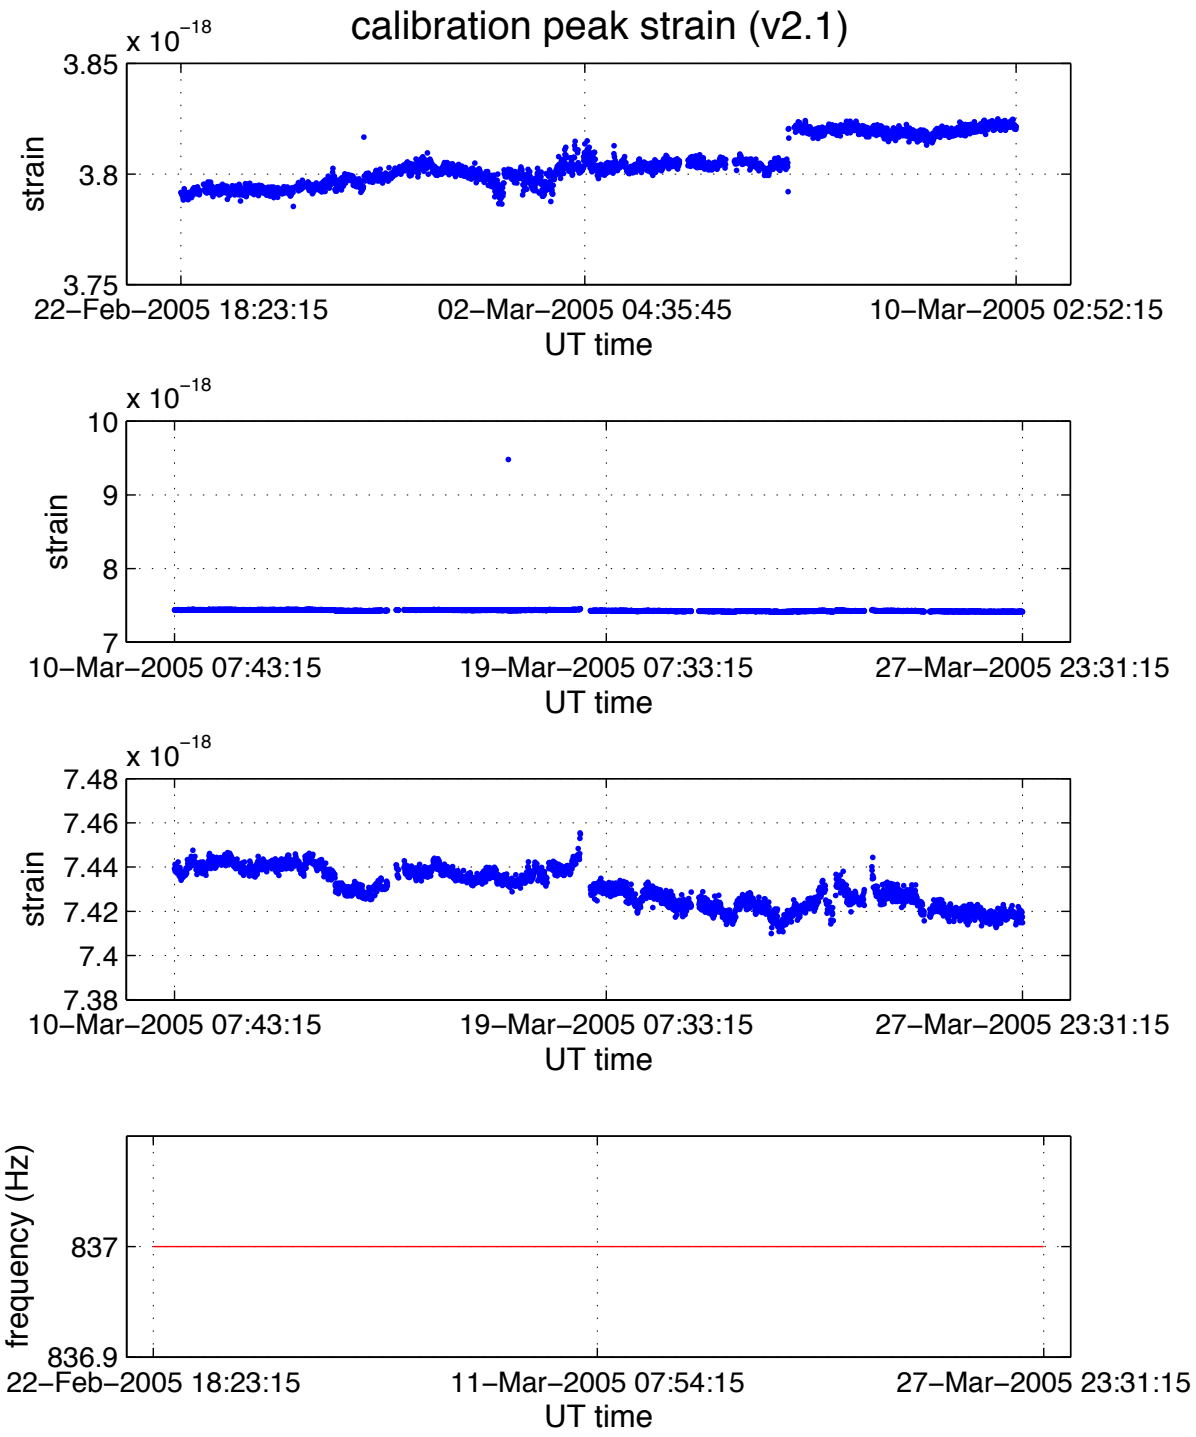

Figure 15: Stability of the calibration line in strain during S4 (`https://sam.phys.lsu.edu/Data_Analysis/309`). There are two epochs corresponding to different DC voltages on the calibrator. The 837 Hz AC voltage remained constant. The first panel shows the epoch with VDC=-4V, the second panel shows the epoch with VDC = -6V. We see a single large outlier at about 25% greater amplitude than the rest. The third panel zooms in ignoring this outlier to see that the rest of that epoch was stable to better than 1%. The last panel shows the frequency of the line. The spectra were calculated at 1/24 Hz resolution and the peak of the line fell into the 837.000Hz bin for all of the data.

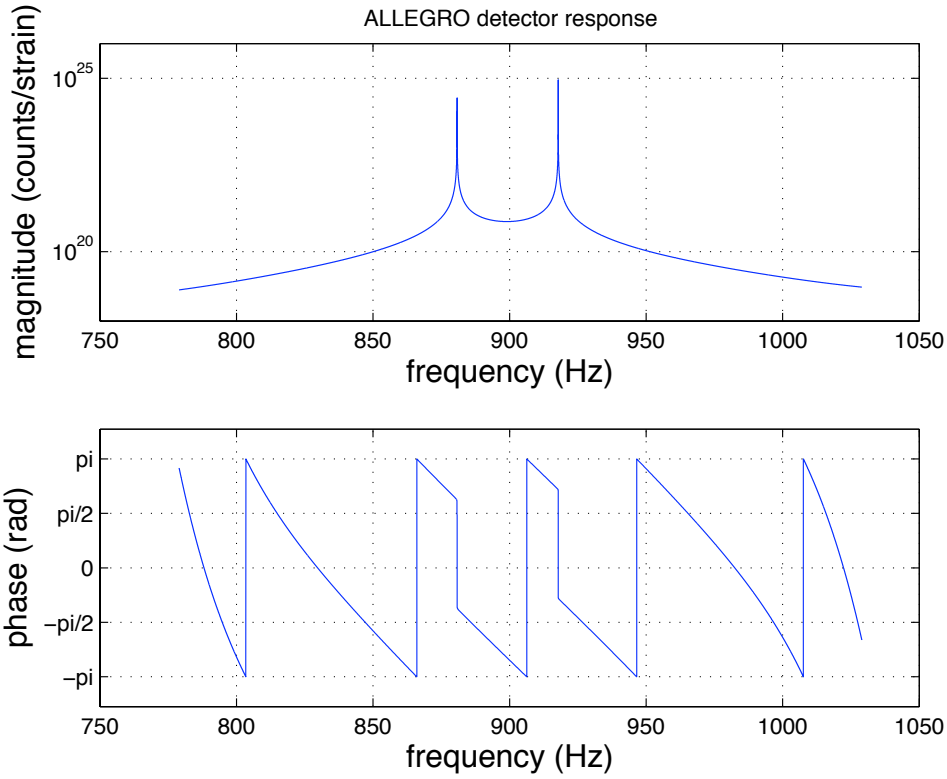

Figure 16: For completeness, we graph the full frequency response function of the ALLEGRO detector – the product of transfer functions $JGKZAD$ in the notation of Sec. 4

Post heterodyne timing Phil Miller's histogram (`https://sam.phys.lsu.edu/Data_Analysis/307`) shows that the variation in 1pps timing over the S4 run is about $\pm 2\mu$s. However, the DAQ timing correction is only made to $\pm 5\mu$s. The residuals shown in Warren's document (`https://sam.phys.lsu.edu/Data_Analysis/275`) are $\pm 10\mu$s. All of these speak to the relative timing. The absolute DAQ timing (`https://sam.phys.lsu.edu/Data_Analysis/276`) uncertainty estimate is $\pm 100\mu$s. The lowpass filter delay is specified to $\pm 50\mu$s. In `https://sam.phys.lsu.edu/Data_Analysis/291` it is concluded that the uncertainty is $\pm 100\mu$s for the filter delay. Thus the dominant terms are the absolute DAQ timing uncertainty and the lowpass filter delay uncertainty. Combining these in quadrature gives $\pm 140\mu$s. We can combine the pre and post heterodyne uncertainties if we specify the band over which the uncertainty holds. For the range 850-960 Hz we get a combined phase uncertainty of $\pm 4°$. The band includes very nearly all of the LLO-ALLEGRO stochastic search optimal filter contribution.

7 Implementation of the transformation

The back-transform procedure is implemented using a set of Matlab routines. These are kept in a CVS respository, **calib_cvs**, under the root **phgrav.phys.lsu.edu:/home/mchugh/cvsroot**. The files which generated the final data used by the stochastic analysis have the tag **ver2p1** – for version 2.1. On the machine **phgrav** at LSU in the directory **/home/monitor/calib_cvs/matlab/cal_2005/** the following was executed to generate the $h(t)$ data files:

```
nohup nice matlab < calibJob.m > calibJob.out &
```

The file **S4params.m** contains most of the parameters that implement the calibration as discussed in this document. We provide here some numbers in terms of the symbols in the overall frequency domain calibration equation $\tilde{h} = \tilde{z}/JGKZAD$ (Eq. 4.15). We also make reference to the Matlab variables in **S4params.m**. These are indicated in **bold**.

We produce z, the complex detector output, by combining the two output channels of the lock-in as recorded by the DAQ. $z = \mathrm{ch2} - i\mathrm{ch1}$ (`https://sam.phys.lsu.edu/Data_Analysis/78`)

$D = 4795.146310$ counts/V – **digitization** – this can be read from the raw data header

$A = a_L e^{i(t_d 2\pi(f-f_r))} e^{-i\phi_L}$ the lock-in amplifier

- $a_L = 1370$ – **lock-in_gain** – for 10 mV sensitivity setting, we also include frequency dependence due to rolloff of the various filters **gfitcoeff.c1, gfitcoeff.c2, gfitcoeff.c3, gfitcoeff.c4**
- $t_d = 12.14$ ms– **tdelay**

- $\phi_L = 15.0°$ – **phase_shift** – there is also a time dependent piece coming from the drift of the reference oscillator – **osc_phase_drift**
- there are also higher order terms **phfitcoeff.c2** - **phfitcoeff.c7**

$Z = 5.2083 \times 10^5$ V/A – **Zsq** – value for the SQUID controller on the 50 ϕ_0 full scale setting.

$K = 196700$ A/m – obtained with **trans_current.value** divided by **trans_eff_gap** – this is for 14 A of stored DC current in the transducer.

$G = \alpha(\frac{1}{(f^2 - f_m^2 - if_m f/Q_m)} - \frac{1}{(f^2 - f_p^2 - if_p f/Q_p)})$ $\alpha = 8.2027 \times 10^{-4}\text{kg}^{-1}$ $f_m = 880.773$ Hz $Q_m = 1.8 \times 10^5$ $f_p = 917.814$ Hz $Q_p = 5.9 \times 10^5$ these are only nominal mode frequencies, as they are tracked throughout the data set, the mode Q 's are given by – **Qmin** and **Qplus**

$J = -4MLf^2$ $M = 2296$ kg – physical mass of the bar $L = 3.0$ m – physical length of the bar

As implemented we have **bar_trans_gain** $= 4ML\alpha$

Another consideration is how the data are subdivided before the inverse response is applied in the frequency domain. Since the mechanical response of the detector G has narrow resonant peaks, long duration stretches of data are used to get the necessary resolution. In practice the raw data are taken in 30 minute stretches. First a Hann windowing function is applied to the data in the time domain to reduce the spectral leakage, then the data are FFT'd. Dividing by the response function as in Eq. 4.15 gives the Fourier coefficients of the effective strain applied to the bar. An inverse FFT brings us back to the time domain. We then shift 15 minutes to get the next 50% overlapping segment of raw data and repeat the process. These stretches are added together to produce the heterodyned strain time series. This series can now be taken in shorter stretches for further analysis as the sharp spectral features have been removed. Fig. 17 shows the amplitude spectral density for a stretch of data that has undergone the calibration procedure.

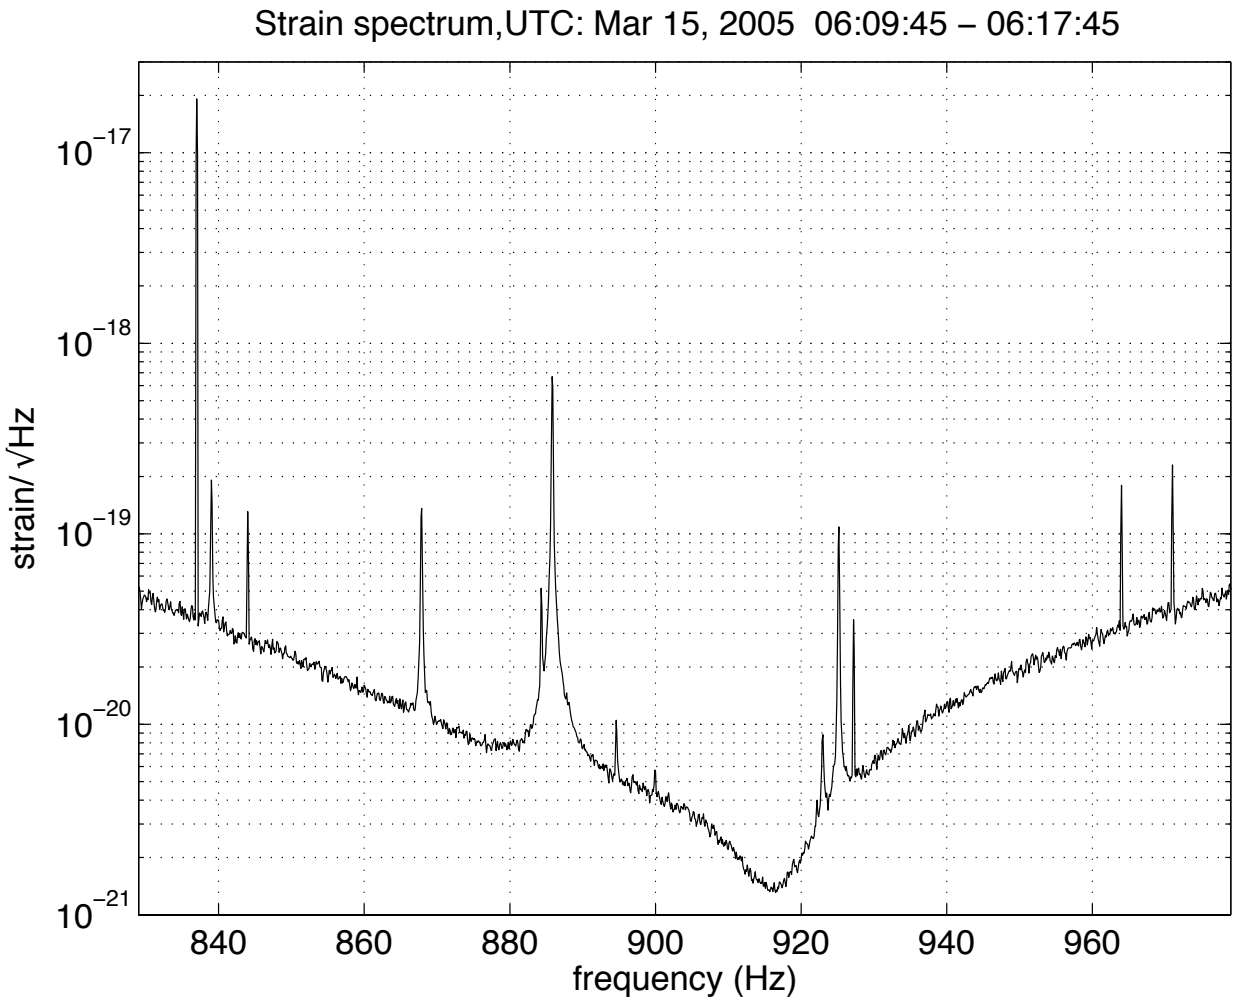

Figure 17: The figure shows the ALLEGRO strain noise curve for operation as during S4. This the same data as in Fig. 13 transformed to the nominal strain variable h. The peak in the noise at the upper mode frequency (917.8 Hz) has changed to a dip, showing that the peak in the response function at this frequency has overpowered the peak in the noise, and so made a dip here. It was a surprise that this did not *also* happen at the lower mode frequency (880.8 Hz) like in the past; the reason was greatly increased mechanical friction (aka much poorer Q) for the lower mode, due to a "touch" between the small mass resonator and the nearby sensing coil.

8 Conclusions

We have outlined the procedure developed for calibration of the ALLEGRO detector. Details of the calibration measurements, as well as more complete examination of the transformed data is forthcoming is a longer publication. The data obtained are being used in a coherent search for a stochastic background of gravitational waves.

9 acknowledgments

We would like to thank our colleagues in the LSC stochastic sources group. In particular Sukanta Bose, Harry Ward and John Whelan have provided helpful input as this calibration procedure has been developed and implemented. This work was supported by the National Science Foundation under grants PHY-0355372 (Loyola) and PHY-9970742 (LSU),

A The tuning curve

Shortly after each cool-down of the antenna, the "tuning curve" was measured. The amount of persistent DC current (I_{DC}) trapped in the superconducting pickup coils of the transducer was systematically varied, and the resulting normal mode frequencies were measured. The procedure required adding extra helium gas around the bar (for extra cooling power), and included a special sequence of changing currents in the "heat switchs" used for steering current through different parts of the circuit. These measurements for the S4 run were done on Feb 13, 2004, see `https://sam.phys.lsu.edu/ALLEGRO/781` .

The magnetic pressure of the persistent current causes extra restoring force on the resonator mass, and so adds an extra "magnetic spring" between it and the end of the bar which should be strictly proportional to I_{DC}^2. To model its effect, it is sufficient to solve for the normal mode frequencies of the model shown in Fig. 18, using an augmented spring constant k'_{34}, given by:

$$k'_{34} = k_{34} + \beta I_{DC}^2 \tag{A.1}$$

where the magnetic spring contribution is parameterized by the coefficient β.

In this model, all masses are determined by theory (for the bar) and material density and geometry. The only free parameters are the static spring constants and the beta factor. Apparently, a good fit of this model to the data is shown in Fig. 19. Shown in the upper two panels are the frequencies versus current, and in the last panel, the quadrature sum vs the quadrature difference of the data in the earlier panels. It can be seen that the fit is good, *but* this particular fit required assuming m_4 was 80% larger than its physical value, which is far greater than the estimated error of a few percent. A similarly good fit could be found by decreasing the bar effective masses $m_2 = m_3$ by the same factor of 1.8, but again, this is not

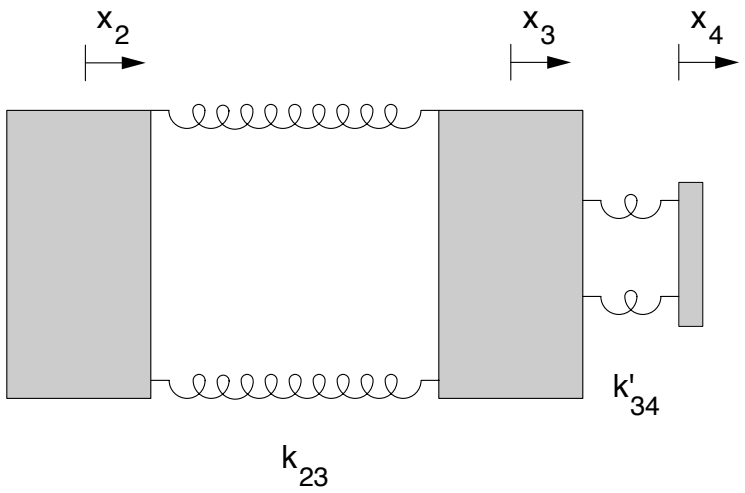

Figure 18: The lumped mass model expected to explain the tuning curve. The resonator spring constant is shown with a prime k'_{34} to indicate it has been augmented by a "magnetic" contribution from the persistent current.

physically possible. (We were able to verify our calculation above of the bar's effective mass, using a simple serial spring-mass model.) Exploring the parameter space quickly convinced us that there was no fit consistent with all the measurements.

Vol Moody, of the University of Maryland, who helped build this transducer, has suggested an alternative model which simply fits the data. (see `https://sam.phys.lsu.edu/Data_Analysis/30` and `https://sam.phys.lsu.edu/Data_Analysis/23`). He assumed there was significant compliance in the mounting flange for the transducer, shown in Fig. 20. The outer case of the transducer is no longer absorbed into the effective bar mass m_3, but is a separate mass $m_c = 11$kg. The resonator mass $m_4 = 0.632$ kg, as determined by its geometry and the density of niobium, both at room temperature.

Varying only the spring constants and β produces a good fit, which is unique, as shown in Fig. 21. The spring constants, parameterized in terms of 'bare' oscillator frequencies, are

$$k_{23} = m_3(2\pi 917.2\text{Hz})^2 \tag{A.2}$$

$$k_{3c} = m_c(2\pi 1800\text{Hz})^2 \tag{A.3}$$

$$k_{c4} = m_4(2\pi 860\text{Hz})^2 \tag{A.4}$$

and the magnetic parameter β was found to be

$$\beta = (\frac{320}{(\text{A}\cdot\text{s})^2})m_4(2\pi)^2 \tag{A.5}$$

in good agreement with the value $\beta = (\frac{315}{(\text{A}\cdot\text{s})^2})m_4(2\pi)^2$ found earlier at Maryland.

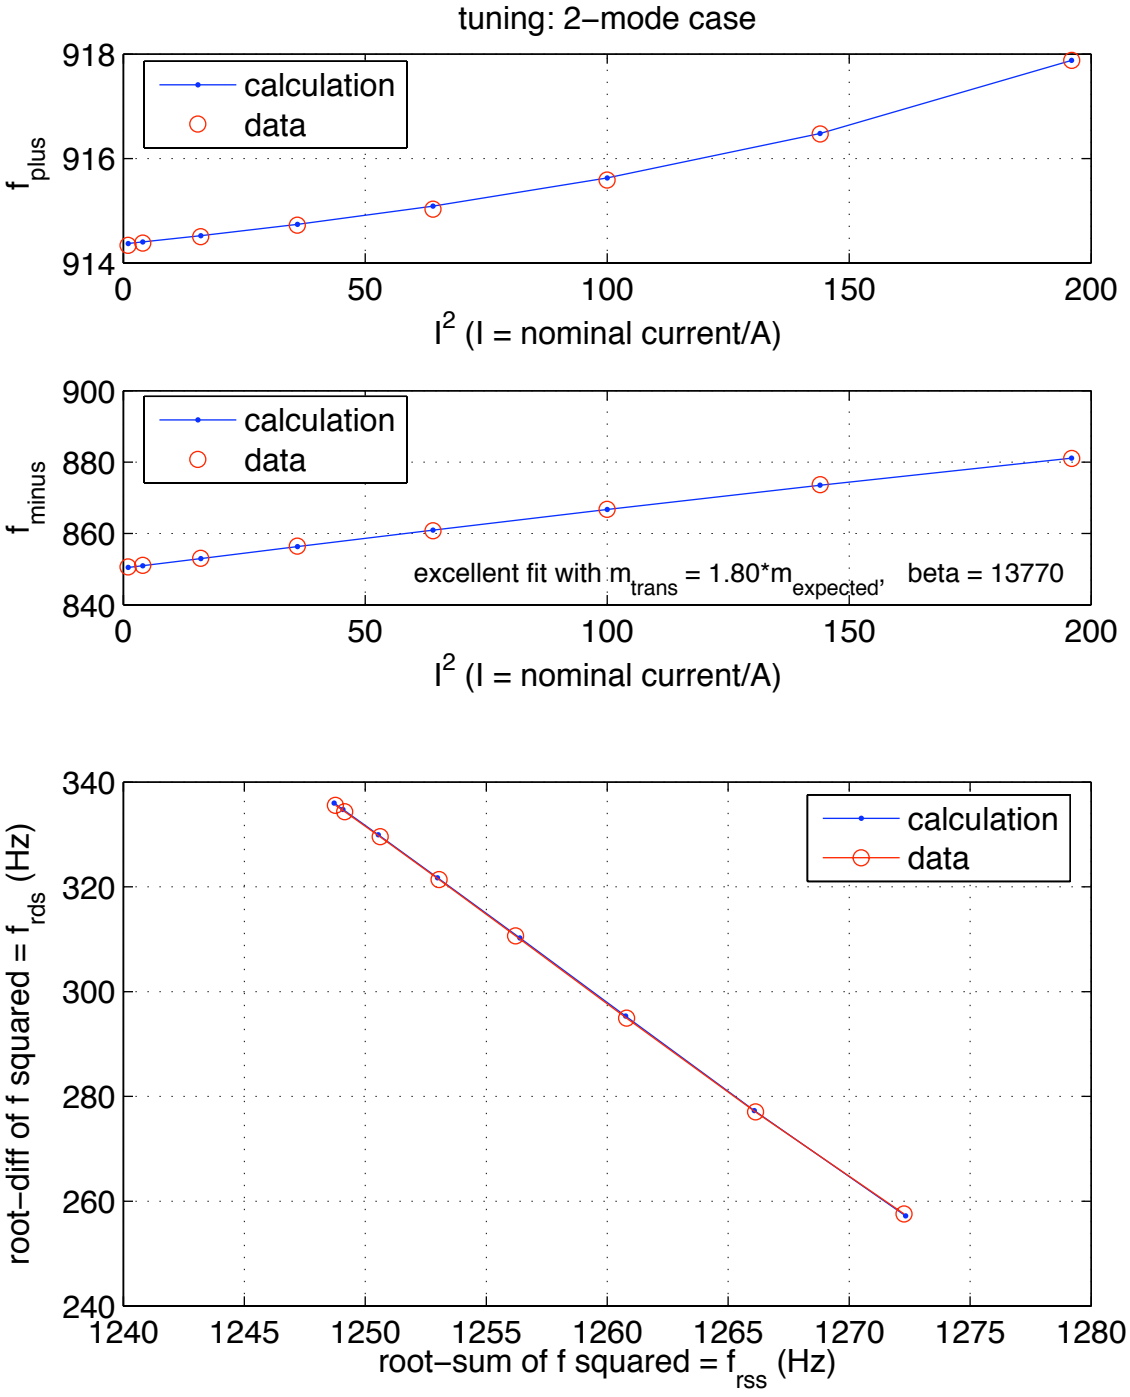

Figure 19: Here is a fit to the data for the model in Fig. 18 which appears to be fine, *but* it is unphysical, because it required allowing the transducer mass m_4 to be a free parameter, arriving a a value 1.80 times larger than its actual value. No manipulation of the real free parameters produced a reasonable fit.

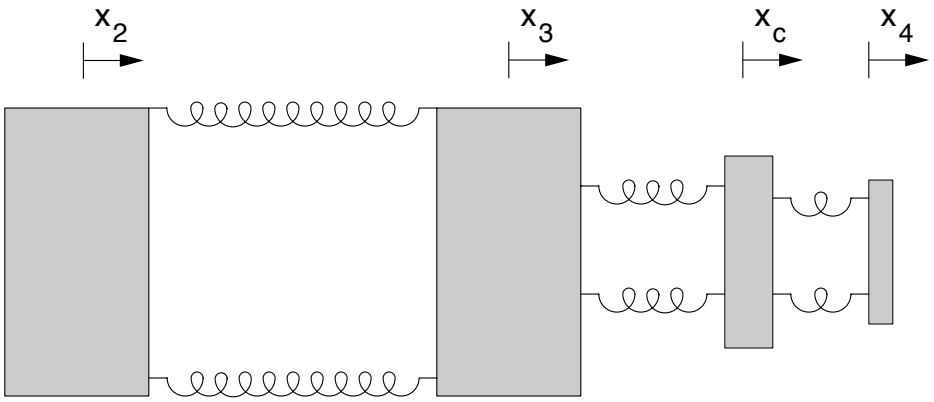

Figure 20: The lumped mass model used to explain the tuning curve, or "Moody's model". The mass of the outer case of the transducer, shown as m_c, is now considered as a separated from the bar, m_3, and now has a finite compliance in its connection to the bar.

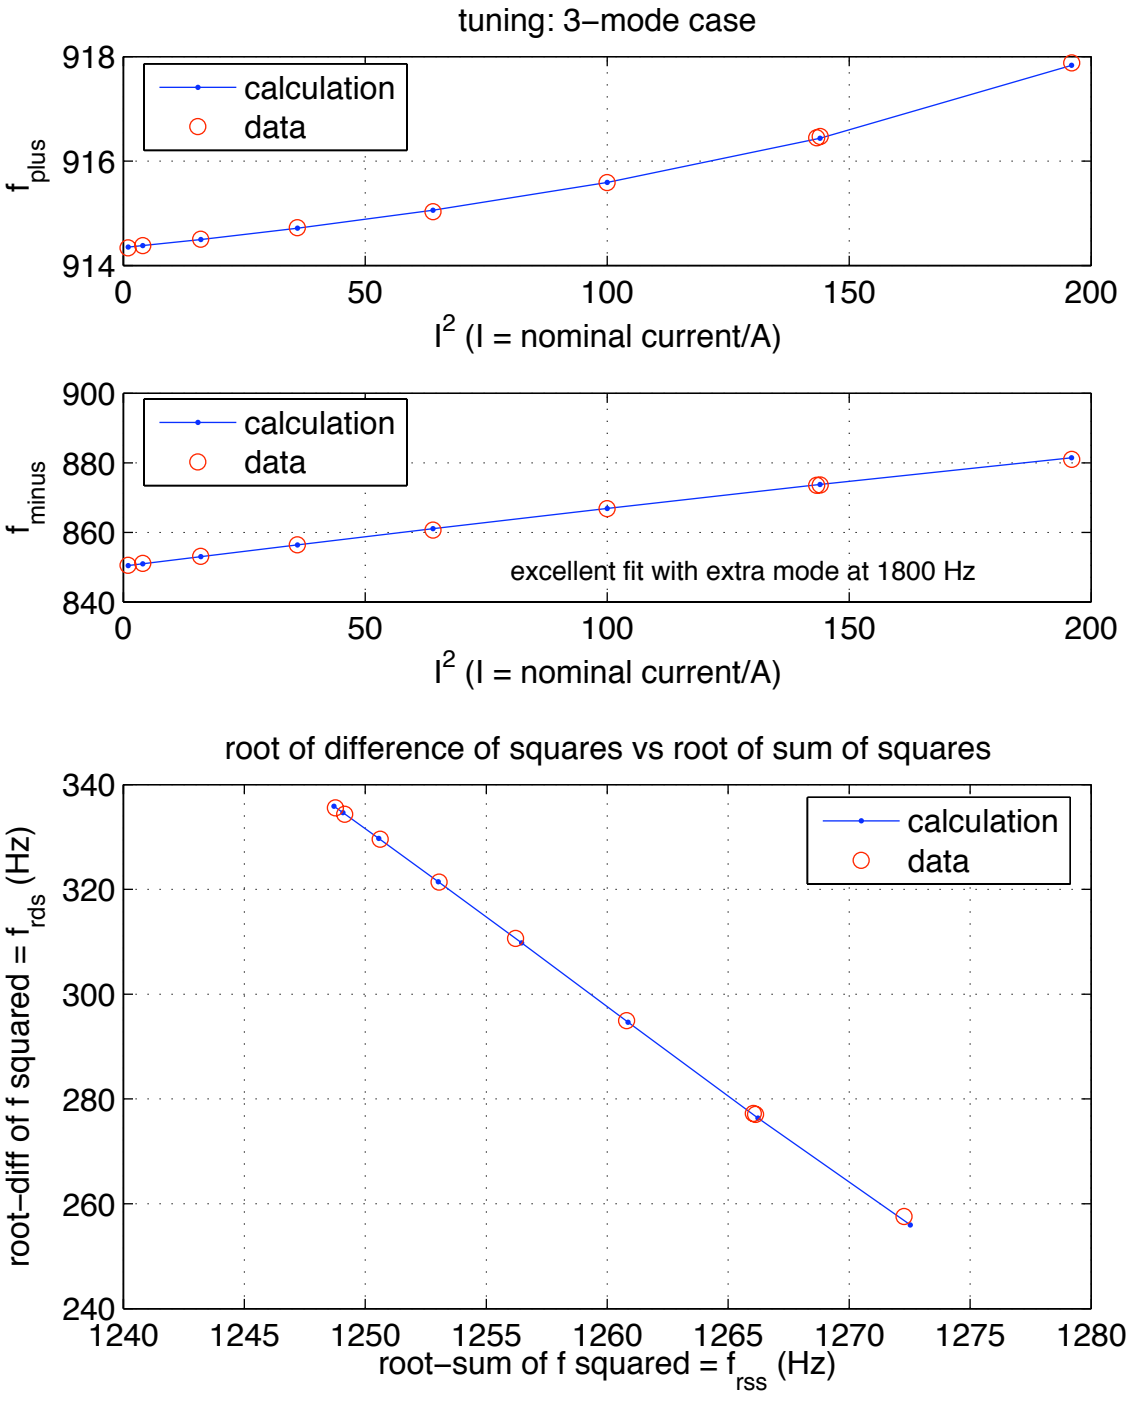

Figure 21: Here is a fit to the data using Moody's model. It assumes that m_c is 11 kg, and finds a good fit for a compliance that should resonate with m_c at 1800 Hz.

B The normal mode decomposition

We restrict ourselves to motion in 1-dimension (≡x) for N masses, so each mass m_i has a displacement $x_i(t)$ and is acted upon by an external force $f_i(t)$ in the x-direction. Then the position and force column vectors $\bar{x}$ and $\bar{f}$ are defined by

$$\bar{x}(t) \equiv \begin{pmatrix} x_1(t) \\ \vdots \\ x_N(t) \end{pmatrix}, \quad \bar{f}(t) \equiv \begin{pmatrix} f_1(t) \\ \vdots \\ f_N(t) \end{pmatrix} \tag{B.1}$$

We define the (diagonal) mass matrix $\mathbf{M}$ and its diagonal "root" $\mathbf{\Gamma}$ by

$$\mathbf{M} \equiv \begin{pmatrix} m_1 & & \\ & \ddots & \\ & & m_N \end{pmatrix} \quad \mathbf{\Gamma} \equiv \begin{pmatrix} m_1^{-\frac{1}{2}} & & \\ & \ddots & \\ & & m_N^{-\frac{1}{2}} \end{pmatrix} \tag{B.2}$$

so

$$\mathbf{\Gamma}^T \cdot \mathbf{M} \cdot \mathbf{\Gamma} \equiv \mathbf{1} \tag{B.3}$$

Then we write out the complete equations of motion, and put all the spring force terms into a symmetric elastic matrix $\mathbf{K}$, so the Newtonian equations of motion become:

$$\mathbf{M} \cdot \ddot{\bar{x}}(t) = -\mathbf{K} \cdot \bar{x}(t) + \bar{f}(t) \tag{B.4}$$

The mass matrix can turned into the unit matrix by transforming to mass-weighted variables $\bar{y}$, defined by

$$\bar{x} \equiv \mathbf{\Gamma} \cdot \bar{y} \tag{B.5}$$

so the equation of motion are now

$$\mathbf{1} \cdot \ddot{\bar{y}}(t) + \mathbf{L} \cdot \bar{y}(t) = \mathbf{\Gamma}^T \cdot \bar{f}(t) \tag{B.6}$$

where

$$\mathbf{L} \equiv \mathbf{\Gamma}^T \cdot \mathbf{K} \cdot \mathbf{\Gamma} \tag{B.7}$$

is the mass weighted elastic matrix, a symmetric matrix.

Then we can simplify further by transforming the position vector to $\bar{z}$ with matrix $\mathbf{V}$

$$\bar{y} \equiv \mathbf{V} \cdot \bar{z} \tag{B.8}$$

and demand that $\mathbf{V}$ diagonalize $\mathbf{L}$, or

$$\mathbf{V}^T \cdot \mathbf{L} \cdot \mathbf{V} = \mathbf{\Lambda} \equiv \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_N \end{pmatrix} \tag{B.9}$$

in other words, we have chosen $\mathbf{V}$ so that the $\bar{z}$ are the normal coordinates, ones which obey uncoupled equations of motion

$$\mathbf{1} \cdot \ddot{\bar{z}}(t) + \mathbf{\Lambda} \cdot \bar{z}(t) = \mathbf{V}^T \cdot \mathbf{\Gamma}^T \cdot \bar{f}(t) \tag{B.10}$$

Most of the information we seek is found in the eigenvector matrix $\mathbf{V}$ with elements V_{ia}, whose 1st index, i, corresponds to the coordinate x_i (of mass m_i). Its second index, a, corresponds the normal coordinate z_a, with eigenvalue $\lambda_a \equiv \omega_a^2$, and mode frequency $f_a \equiv \omega_a/(2\pi)$. $\mathbf{V}$ is *not* symmetric; but it is a set of orthonormal column vectors, so

$$\mathbf{V}^T \cdot \mathbf{V} = \mathbf{1} \tag{B.11}$$

When indices are explicitly displayed, $\mathbf{V} => V_{ia}$ implies $\mathbf{V}^T => V_{ai}$. In explicit index notation, Eq. B.10 becomes

$$\ddot{z}_a(t) + \omega_a^2 z_a(t) = \sum_{i=1}^{N} V_{ia} m_i^{-\frac{1}{2}} f_i(t) \tag{B.12}$$

which can be solved by fourier analysis. Then the general solution, in the fourier domain, is found by transforming back to the displacement $\tilde{x}_j(\omega)$ of each mass m_j

$$\tilde{x}_j(\omega) = \sum_{i=1}^{N} \left[\sum_{a=1}^{N} m_j^{-\frac{1}{2}} V_{ja} \frac{1}{(-\omega^2 + \omega_a^2)} V_{ia} m_i^{-\frac{1}{2}} \right] \tilde{f}_i(\omega) \tag{B.13}$$

where the Green function matrix element for the system is contained within the square brackets, with output index j and input index i. Each matrix element contains a sum over mode index a. Notice the reversed order of indices in the two terms of the eigenvector matrix V.

As a simple example, let us calculate the "mode mass" of the "two mass" model for a bar. The model is *two* point masses m, connected by a spring k. The elastic matrix is

$$\mathbf{K} = \begin{pmatrix} k & -k \\ -k & k \end{pmatrix}, \tag{B.14}$$

so the weighted elastic matrix factors into a parameter and a matrix of integers,

$$\mathbf{L} = \begin{pmatrix} k/m & -k/m \\ -k/m & k/m \end{pmatrix} = (k/m) \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}. \tag{B.15}$$

Then we find the eigenvectors $\mathbf{V}_1$ and eigenvalues $\mathbf{\Lambda}_1$ of the integer matrix to be

$$\mathbf{V}_1 = \begin{pmatrix} -1/\sqrt{(2)} & 1/\sqrt{(2)} \\ -1/\sqrt{(2)} & -1/\sqrt{(2)} \end{pmatrix} \quad \mathbf{\Lambda}_1 = \begin{pmatrix} 0 & 0 \\ 0 & 2 \end{pmatrix} \tag{B.16}$$

Then $\mathbf{L}$ has the *same* eigenvector matrix $\mathbf{V} = \mathbf{V_1}$, and the eigenvalues are multiplied by the factor (k/m).

So, for the first mode, $a = 1$, the resonant frequency ω_a is zero, and by substitution

$$\begin{pmatrix} \tilde{x}_1(\omega) \\ \tilde{x}_2(\omega) \end{pmatrix} = \left(\frac{(m)^{-1}}{(-\omega^2 + 0)} \right) \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix} \begin{pmatrix} \tilde{f}_1(\omega) \\ \tilde{f}_2(\omega) \end{pmatrix} \tag{B.17}$$

which reduces to

$$\tilde{x}_1(\omega) = \tilde{x}_2(\omega) = \left(\frac{(2m)^{-1}}{(-\omega^2)} \right) (\tilde{f}_1(\omega) + \tilde{f}_2(\omega)) \tag{B.18}$$

which says that the two masses travel as one, responding like a single particle with mass $2m$, accelerated by the sum of the forces.

$$\ddot{x} = \frac{f_1 + f_2}{2m} \tag{B.19}$$

As for the second mode, $a = 2$, its resonant frequency follows from

$$\lambda_2 = (k/m)(2) = \omega_2^2 \tag{B.20}$$

and now, from the second column of V, the response matrix becomes

$$\begin{pmatrix} \tilde{x}_1(\omega) \\ \tilde{x}_2(\omega) \end{pmatrix} = \left(\frac{(m)^{-1}}{(-\omega^2 + \omega_2^2)} \right) \begin{pmatrix} 1/2 & -1/2 \\ -1/2 & 1/2 \end{pmatrix} \begin{pmatrix} \tilde{f}_1(\omega) \\ \tilde{f}_2(\omega) \end{pmatrix} \tag{B.21}$$

Now if we specialize to a single force f_1, acting only on mass m_1, and ask for the response at the driving point, x_1, we get

$$\tilde{x}_1(\omega) = \left(\frac{(2m)^{-1}}{(-\omega^2 + \omega_2^2)} \right) \tilde{f}_1(\omega) \tag{B.22}$$

which is the response of a "harmonic oscillator" with mode mass $2m$. We could have achieved this result easier by simpler means (like adding and subtracting the equations of motion), but they would not generalize to more complicated models.

Compare this to the driving point inertia we find for the face of a bar, with physical mass M, near its first resonance, which we found to be $M/2$. So if we wish to model a bar with *two* masses coupled only to each other, then we need $2m = M/2$, or $m = M/4$.

References

[1] http://gravity.phys.lsu.edu/

[2] Gregory M. Harry, thesis, University of Maryland 1999

[3] Astone P et al 2003 *Phys. Rev.* **D68** 022001

[4] Mauceli E, Geng Z K, Hamilton W O, Johnson W W, Merkowitz S, Morse A, Price B and Solomonson N 1996 *Phys. Rev.* **D54** 1264-1275 (Mauceli E et al 1996 *Preprint* gr-qc/9609058)

[5] Morse A, Hamilton W O, Johnson W W, Mauceli E and McHugh M P 1999 *Phys. Rev.* **D59** 062002

[6] Allen B and Romano J D 1999 *Phys. Rev.* **D59** 102001; (Allen B and Romano J D 1997 *Preprint* gr-qc/9710117)

[7] Kolb E W and Turner M S 1990 *The Early Universe* (Addison-Wesley, 1990)

[8] Whelan J T, Daw E, Heng I S, McHugh M P and Lazzarini A 2003 *Class.Quant.Grav.* **20** S689-S695 (Whelan J T et al 2003 *Preprint* gr-qc/0308043)

[9] Finn L S and Lazzarini A 2001 *Phys. Rev.* **D64** 082002; (Finn L S and Lazzarini A 2001 *Preprint* gr-qc/0104040)

[10] http://www.ligo.caltech.edu

[11] J. T. Whelan, S. Bose, J. Hanson, I. S. Heng, W. W. Johnson, M. P. McHugh and P. Zhang 2005 *Class.Quant.Grav.* **22** S1087-S1096

[12] Carroll Andrew Morse, thesis, Louisiana State University 1999

[13] M. P. McHugh, W. W. Johnson, W. O. Hamilton, J. Hanson, I. S. Heng, D. McNesse, P. Miller, D. Nettles, J. Weaver and P. Zhang 2005 *Class.Quant.Grav.* **22** S965-S973