

Chapter 3: BOOMERANG Prototype

3.0 Overview

This chapter details the construction of a prototype BOOMERANG spectrometer optimized for detection of NMR from solid and liquid samples contained in a 3 mm diameter sample volume. Figure 3.1 shows a block diagram of the prototype. The heart of the spectrometer is the magnet assembly, which is




Figure 3.1. Block diagram of the BOOMERANG prototype.

described in section 3.1. This assembly provides a homogeneous field for a sample that is placed inside a solenoid, which is situated underneath a sensor magnet (not shown) inside the assembly. The sensor magnet is bound to a single-crystal silicon suspension (also not shown), and this combination forms a mechanical oscillator, which is described in section 3.2. Picometer-scale motions of the oscillator are monitored with the fiber-optic interferometer system described in section 3.3. The oscillator's resonance frequency is typically between 400 and 600 Hz. In order to suppress ambient acoustic noise, the magnet assembly is enclosed in a vacuum bell jar, which is pumped down to below 10^{-5} Torr. The vacuum also serves to suppress viscous damping of the oscillator.

Typically, the effects of viscous damping are observed to become negligible relative to eddy-current damping when the pressure is below about 10^{-3} Torr at room temperature. Under these conditions, the mean free path of molecules in the air is about 76 mm, which is far larger than the dimensions of the sensor magnet. Thus, a continuum model of the rarified atmosphere, which would predict that the air viscosity is independent of pressure¹, does not apply. We observe a modest increase in the ring-down time of the oscillator up to as much as 80 ms as the pressure is reduced below ambient pressure.

The optical signal from the fiber-optic position sensor is filtered, amplified, and digitized (section 3.4). Rf pulses to reorient the sample's magnetization, including the frequency-modulated pulses used to efficiently invert the magnetization during detection, are digitally synthesized, mixed up to the proton or fluorine Larmor frequency in the 25-30 MHz range, amplified, and delivered to an rf




Figure 3.2 BOOMERANG prototype magnet assembly. The drawing is to scale, and magnetic elements, drawn in solid lines, are cylindrically symmetric about an axis that is vertical in the figure. All the magnets are magnetized along their common axis of symmetry. Square, brass brackets to hold the magnets are drawn in dotted lines. Magnetostatic calculations to specify exact magnet dimensions and spacings, as well as designs of bracket and alignment hardware, were performed by Lou Madsen².

coil (the solenoid containing the sample) inside the magnet assembly as described in section 3.5. Section 3.6 treats cyclic inversion, which is used to drive the detector oscillator in proportion to the sample's magnetization. The chapter concludes with a detailed assessment of noise sources (section 3.7).

3.1 Magnet Assembly

The defining feature of the BOOMERANG magnetic resonance spectrometer is the magnet assembly, which is designed specifically for field homogeneity. Figure 3.2 shows a schematic of the magnet assembly, which is to-scale. The aluminum base and supports used to hold the magnets in alignment are left out of the diagram for clarity. The "source" magnets (Magnetic Component Engineering, Inc., model N40) are made of neodymium iron boron (NdFeB) magnetized axially to a remanent

magnetization of $1.29 \text{ T}/\mu_0$. They are right circular cylinders 2.54 cm high and 5.08 cm in diameter. Both have 1.6 mm holes drilled down the symmetry axis, the upper magnet to accommodate an optical fiber, the lower magnet to preserve reflection symmetry of the magnet assembly.

These source magnets magnetize the four internal magnets, including the sensor magnet. The internal magnets are machined from mu metal (Carpenter Technology Corp. "HyMu 80" alloy). In the field of the source magnets, the internal magnets acquire a magnetization of $0.665 \text{ T}/\mu_0$. This value equalizes the normal component of the magnetic induction B across the boundary between the high-permeability mu metal and the air gap for the known geometry of the magnet assembly. This estimate, which falls between the value calculated for an infinitely wide flat disk in the field of the source magnets ($0.605 \text{ T}/\mu_0$) and the material's saturation magnetization³ ($0.75 \text{ T}/\mu_0$) is the result of an iterative magnetostatic self-consistency calculation.

Figure 3.3 shows a plot of the longitudinal component B_z of the static field along the symmetry axis in the 4.8 mm space below the sensor magnet and above its complement. The distance between the annular magnets is fixed, but the distance between the source magnets can be varied somewhat to provide a coarse shimming capability. This procedure also changes the static field, and we have observed fields across the range 0.59-0.73 T. (The magnetization of the sensor magnet changes accordingly.) The three curves in Figure 3.3 are for three positions of the sensor magnet, which is free not only to vibrate longitudinally, but also to come to rest slightly above or below its encircling annulus as a result of static




Figure 3.3. Calculation of the field at positions along the symmetry axis in the BOOMERANG prototype. When the sensor magnet is perfectly centered in the annulus, odd-order gradients in the field vanish.

forces from the other magnets. By varying the distance between the source magnets with inserted brass shims, it is possible to align the sensor and annulus to within $\sim 50 \mu\text{m}$ by careful visual inspection.

The reflection symmetry of the magnet system makes all odd-order gradients (*e.g.*, $\partial B_z/\partial z$, $\partial^3 B_z/\partial z^3$) vanish. The principal design goal was therefore to null the second-order gradient $\partial^2 B_z/\partial z^2$ as best as possible. As is shown in Figure 3.3, the linear term vanishes only when the sensor magnet is well aligned. Figure 3.4 shows three line shapes calculated for the three positions (perfect alignment and $25 \mu\text{m}$ above and below perfect alignment) of Figure 3.3. These line shapes are calculated by computing the field expected at randomly chosen points in a 3 mm spherical volume, binning the values, and graphing the resulting probability distributions. The simulations show that the total distribution of proton Larmor frequencies over the




Figure 3.4. Distribution of fields within the sample for three positions of the sensor magnet.

sample is on the order of 28 kHz (6.5×10^{-4} T \times 42.6 MHz/T, with full width at half maximum (FWHM) 6 kHz) for a perfectly aligned sensor and that the line width can vary substantially depending on the sensor's alignment. This distribution of Larmor frequencies is compatible with inversion of the entire sample magnetization for the \sim 50 kHz Rabi frequencies we achieved with our rf system (section 3.5).

Since the force coupling a given nuclear moment to the sensor magnet varies as a function of the moment's position within the sample, the distribution of Larmor frequencies shown in Figure 3.4 must be weighted by a "local sensitivity" in order to properly model the NMR line shape. As shown in Figure 3.5, this weighting distorts the expected line shapes considerably more for spectra taken with the sensor magnet displaced from its most symmetric location.

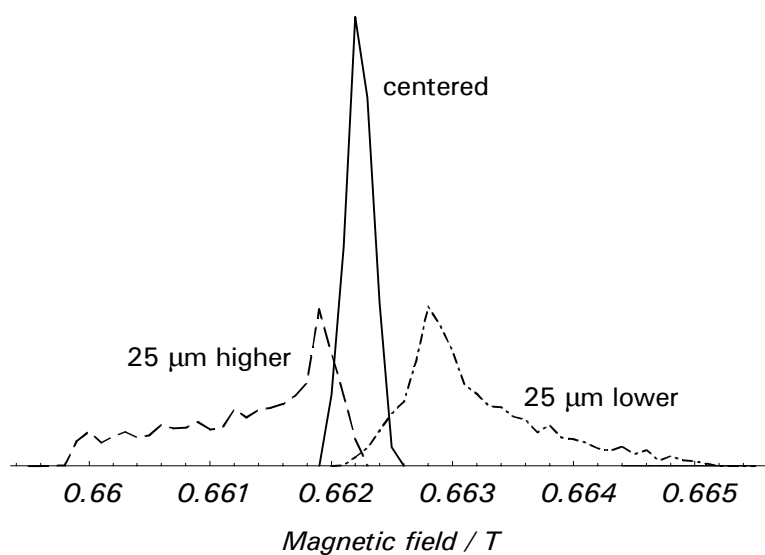


Figure 3.5. Simulated lineshapes for three positions of the sensor magnet adjusted for sensitivity variations in the sample volume. When the sensor is misaligned, extreme values of the field in the NMR line are weighted more significantly.

Figure 3.6 shows a Fourier-transform (FT) ^1H NMR spectrum of water taken with the prototype (see Chapter 4 for details). The line widths observed in ^1H NMR of water were typically in the range 8-50 kHz, which are consistent with the above simulations. Most of the observed line width can be attributed to the field inhomogeneity due to the presence of the gap between the sensor and the annulus. Indeed, simulation of the distribution of Larmor frequencies for the particular geometry in Figure 3.2 with the sensor radius modified so as to have no gap between the sensor and annulus showed a line width of only 0.6 kHz. Detailed optimization of the homogeneity using the approximation of perfectly axial magnetization, as well as more exact finite-element analyses² showed that this line width can be reduced, with careful placement of all the magnets, to well below 1 ppm (28 Hz in the prototype) overall for such no-gap configurations. An important




Figure 3.6. FT-NMR spectrum of water measured with the BOOMERANG prototype. The 17 kHz line width is consistent with the model calculation of figure 3.5.

improvement to next-generation BOOMERANG devices will be an active shimming capability based on movement of the magnets, which could be supplemented by a shim coil set as in ordinary NMR. Even at the demonstrated homogeneity (which is more than three orders of magnitude better than would be the case without using the BOOMERANG concept), the line width is sufficient to allow inversion of the entire sample magnetization. This is all that is really required. In the highest-resolution BOOMERANG designs, spectroscopic evolution can take place with the sample placed at a separate location, which is away from the sensor magnet and its surrounding gap, such a location being separately optimized for strict homogeneity (perhaps with coil-based shimming measures), with the sample subsequently shuttled under the sensor magnet for optimal detection⁴.




Figure 3.7. BOOMERANG prototype oscillator assembly, with surrounding annular and complement magnets, silicon suspension, and NMR coil. a) Side view, 3x scale. The sensor magnet, 3 mm in diameter and 1.5 mm high, is bound to a 0.22 mm thick rectangular silicon suspension. b) Top view, 1x scale.

3.2 Oscillator

Figure 3.7 shows the center of the magnet assembly with its oscillating silicon beam and sensor magnet. The sensor magnet is affixed to the beam so that it is free to oscillate along the symmetry axis. The sensor magnet's counterpart in the lower half of the assembly is glued inside its annulus, its face flush with that of the annulus. The sensor magnet weighs 83.1 mg, and the silicon beam's total mass is 36.7 mg. The total motional mass for this "fixed-fixed" beam configuration⁵ is therefore about $(83.1 + 0.375 \times 0.70 \times 36.7) \text{ mg} = 92.7 \text{ mg}$ accounting for the $\sim 30\%$ of the silicon beam that is fixed to the annulus and not free to oscillate. When placed between the field magnets, the oscillator's resonance frequency was typically between 400–600 Hz. The frequency of a given oscillator varied by tens

of hertz or more depending on the age of the oscillator, the temperature, and the thickness of the brass shims used to separate the source magnets.

The total spring constant, ~ 897 N/m (for the particular case of the above motional mass and a 495 Hz resonance frequency), is the result of two offsetting effects: a positive elastic spring constant due to the restoring force of the silicon beam and a negative magnetic spring constant due to magnetic forces, primarily between the sensor magnet and its encircling annulus. In the absence of the restoring force of the silicon beam, the sensor magnet is at the “top of a hill” in potential energy when positioned at the center of the annulus². A crude estimate of this contribution to the spring constant can be made by estimating the second derivative with respect to longitudinal displacements of this potential energy, with the sensor modeled as a simple dipole in the field of the annulus. The result is -1380 N/m, which indicates an elastic spring constant of about $+2280$ N/m. The negative magnetic contribution to the total spring constant is an important feature of the BOOMERANG method. While both the elastic and magnetic spring constants (and therefore their sum) should scale linearly with the size of the apparatus, variations in aspect ratios of the silicon suspension should allow some control in the adjustment of the mechanical resonance frequency at a given size scale. This will help to maintain a frequency low enough to permit inversion of the sample magnetization with practical rf power (see section 3.5) as BOOMERANG devices are scaled down. Active measures of controlling the spring constant² may further enhance this capability.

The frequency and the ring-down time of the oscillator are experimentally measured by observing the steady-state response to acoustic or magnetic excitation as a function of input frequency. The oscillator is excited by magnetic coupling to the field generated by audio-frequency currents in a nearby excitation coil that is placed temporarily under the magnet. The mechanical response of the oscillator is maximized visually on an oscilloscope while the input frequency is tuned to find the resonance. The line width $\gamma/2\pi$ is obtained by recording those frequencies, one each on the high- and low-frequency sides of the resonance frequency, at which the mechanical response is $1/\sqrt{2}$ times the amplitude at resonance. In cases where the mechanical resonance is particularly sharp, the ring-down time $\tau = 2/\gamma$ of the oscillator can be measured directly by observing the transient response to an impulsive excitation. Typically, the line widths of the best oscillators used in our experiments were in the range 4–6 Hz (ring-down time $\tau = 53$ –80 ms).

The damping of the oscillator is probably dominated by eddy currents induced in both the annulus and sensor magnets by their relative motion. This conclusion is based on the observation that greatly widening the gap (by 0.5 cm or more) between the sensor and the annulus substantially lengthens the ring-down time². Calculations in Appendix C show that these (azimuthal) eddy currents are primarily located very close to the edges of the sensor and annulus. Indeed, longer ring-down times are observed with sensor-ring combinations with slightly rounded edges in the sensor/annulus gap. This suggests that some tradeoff may be made between homogeneity and ring-down time. Another strategy for the reduction of eddy currents that we have used with some success is the introduction of radial

slits in the annulus to interrupt the currents. This approach is motivated by common practice in transformer design, in which laminated transformer cores are used to mitigate similar eddy current losses.

The calculation of eddy current damping in Appendix C assumes uniform axial magnetization in both the sensor and the annulus, and it underestimates somewhat the eddy current damping that is actually observed. Radial components in the magnetization would make the calculated eddy currents larger. In magnetic fields that are substantially larger than the ~ 0.66 T field of the prototype, the magnetization of sensor and annulus may more strongly conform to an axial orientation, which may help to reduce eddy current damping.

Another significant issue regarding the oscillator is the drift of its resonance frequency between iterations of the experiment. This is probably due to heating of the oscillator by the applied rf current. The drift over several shots of the experiment can be seen in the plot of Figure 3.8. This density plot records the Fourier transform of a time series, a record of the oscillator's trajectory during driving by magnetization modulated with rf at a frequency (here 442 Hz) that is fixed shot-to-shot. The dark band is the frequency range over which the Brownian motion of the oscillator is strongest, which is within ~ 1 line width of the resonance frequency. This drift can be compensated for in practice by including several "dummy" applications of rf, which brings the oscillator to a steady-state temperature and frequency.




Figure 3.8. Drift over time of the oscillator frequency. In this density plot, intensity in the Fourier transform of a mechanical transient is indicated by color for 150 iterations of the experiment. The broad band, which moves by ~ 10 Hz over the course of the experiment, is the ~ 7 -Hz-wide signature of the Brownian-motion noise. The narrow, dark feature at 442 Hz is the NMR signal that results from rf frequency modulation during oscillator driving.

Of more concern is the drift of the oscillator *during* a given shot. This can cause phase variations in the signal when the time between the beginning of heat deposition from rf pulses to the onset of the detection period varies from shot to shot in the experiment. Frequency variations of the oscillator were observed in real time by applying an audio-frequency excitation to the oscillator at a fixed frequency slightly (~ 3 Hz) off resonance and observing the phase shift on an oscilloscope as the natural resonance frequency changed during application of the rf current. The

frequency changed by ~ 4 Hz during application of rf currents that were typical of the experiments reported in Chapter 4. This is a very substantial fraction of the natural line width of the oscillator, and it was observed to cause significant aberrations in the phases of time-domain NMR signals.




Figure 3.9. ^{19}F Nutation time series showing phase aberrations. The sample is neat fluoroacetonitrile. The signal should show a single decaying cosine.

Figure 3.9 shows the result of a particularly bad nutation experiment. In this figure, deviations from the expected decaying cosine are far above the predicted (section 3.7) and observed base noise level. Ways to account for this in the experiment are the subject of ongoing design efforts. A promising approach is to concurrently excite the oscillator slightly off resonance and observe the phase shift of this signal in real time. This signal can be used to estimate the instantaneous resonance frequency, which can subsequently be used in a fitting procedure. Alternatively, this estimate, or a temperature measurement, can be used in a thermal feedback scheme to stabilize the oscillator's frequency in real time.

3.3 Fiber-Optic Interferometer

The picometer-scale motions of the sensor magnet are monitored with a fiber-optic interferometer⁶, which is shown schematically in Figure 3.10. Laser light at 780 nm from a pigtailed multimode diode laser (Sharp model LT023MD, from OZ Optics, Canada) is launched into one arm of a 2x2 fiber coupler (Gould Fiber Optics,

Inc.) through a single-mode optical fiber (5 μm core, 125 μm cladding). The laser is driven by a battery-powered current source for minimum intensity noise. The light is split 50/50 into the coupler's two output arms. One arm goes into a light dump to attenuate destabilizing backscattered light that would otherwise return to the laser arm. Dumping the light in this arm also reduces an undesirable dc offset in the detected photocurrent. The light dump is simply a piece of black felt affixed to the end of the fiber and soaked with pump oil, which serves as an index-matching fluid to further reduce reflections off the fiber end. The light dump is simply a piece of black felt affixed to the end of the fiber and soaked with pump oil, which serves as an index-matching fluid to further reduce reflections off the fiber end.

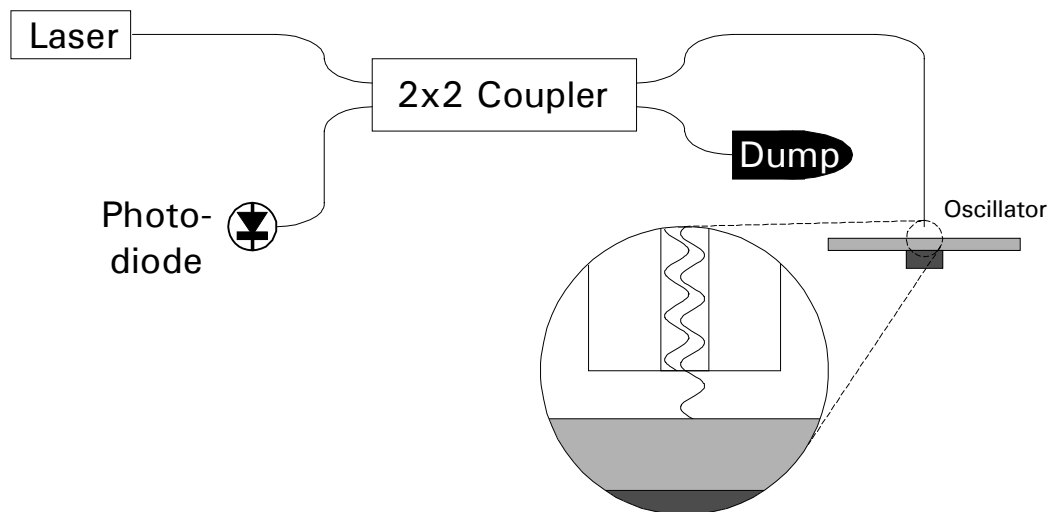


Figure 3.10. Fiber-optic interferometer. A laser diode launches light into one arm of a 2x2 fiber coupler, where it is split between two single-mode fibers. Half of this light goes to a light dump to prevent unwanted reflections. The other half goes into the magnet array through a narrow hole drilled through the center of one NdFeB pole magnet. The end of this fiber is cleaved, and it is brought to within a few microns of the polished surface of the silicon oscillator. Reflections from the glass-air interface and from the polished surface scatter back into the fiber and travel back through the coupler to a photodiode. The displacement of the oscillator is registered as a variation of intensity of the light incident on the photodiode due to interference of these two scattered beams.

The light in the coupler's interferometer arm goes through a vacuum-wall feedthrough, into the vacuum bell jar, and through the hole drilled along the symmetry axis of the upper field magnet. The fiber's face is cleaved flat and

brought to within a few microns of the reflecting surface at the back of the silicon oscillator. The light is reflected from two surfaces — one, the glass-air interface at the fiber end, and the other, the silicon oscillator. If the distance between these two reflecting surfaces is smaller than the coherence length of the laser, then the light reflected back down the fiber toward the coupler is subject to modulation by the interference between these two reflected waves, which depends sinusoidally on the distance between the fiber face and the oscillator. This backscattered light goes back through the coupler, and it is again split 50/50. One half of this light is incident on a photodiode (Seastar model CP-120-20), which is fc-coupled to the fiber. The photodiode is reverse-biased with a 9V battery to improve response and linearity, and the photocurrent is amplified by a current (transimpedance) amplifier (Princeton Applied Research, Model 181). The resulting output voltage is subsequently filtered, amplified, and digitized (section 3.4).

The interferometer-arm fiber is affixed above the source magnet to a clamp-and-spring assembly, which allows the distance between the fiber end and the silicon oscillator to be regulated by applying a voltage to a piezoelectric stack under the clamp. The nominal distance is set to the center of a sinusoidal fringe, where the photocurrent varies linearly with small displacements. Since it is important to maintain this linearity throughout the duration of the experiment, a feedback circuit is used to keep the fiber face near the fringe center². The fringe visibility of the interferometer,

$$V = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \quad (3.1)$$

where I_{\max} and I_{\min} are the maximum and minimum photocurrents, is typically 45-75%.

The sensitivity of the fiber-optic interferometer system is sufficient to observe the Brownian motion of the oscillator at room temperature, and so Brownian motion is the predominant noise source for room-temperature

measurements. Figure 3.11 shows the observed noise spectral density over a 90 Hz range that includes the oscillator's resonance. This resonance peak is fit to a Lorentzian line at

$\omega/2\pi = 496$ Hz with full-width at half-maximum $\gamma/2\pi = 5.0$ Hz. The silicon element of this particular oscillator was part of a structure more

complicated than a single fixed-fixed

beam, and so the motional mass is difficult to estimate from elementary solid

mechanics. The noise spectral density shown is consistent with a motional mass of

139 mg, which is somewhat larger than the 92.7 mg estimated for the simple fixed-

fixed beam. Also included in the fit is a frequency-independent noise floor, which is

about $0.038 \text{ pm}^2/\text{Hz}$ ($0.195 \text{ pm}/\sqrt{\text{Hz}}$). This exceeds estimates of white noise from

the sum of several noise sources quantified in section 3.7 by a factor of four in

power and two in amplitude. This excess noise may be due to intensity noise in the

laser or possibly residual acoustic noise.




Figure 3.11. Displacement noise spectral density. The peak at the oscillator frequency is due to Brownian motion. The frequency-independent noise floor is somewhat larger than the prediction based on shot noise, Johnson noise in the photocurrent amp, etc., calculated in section 3.7.

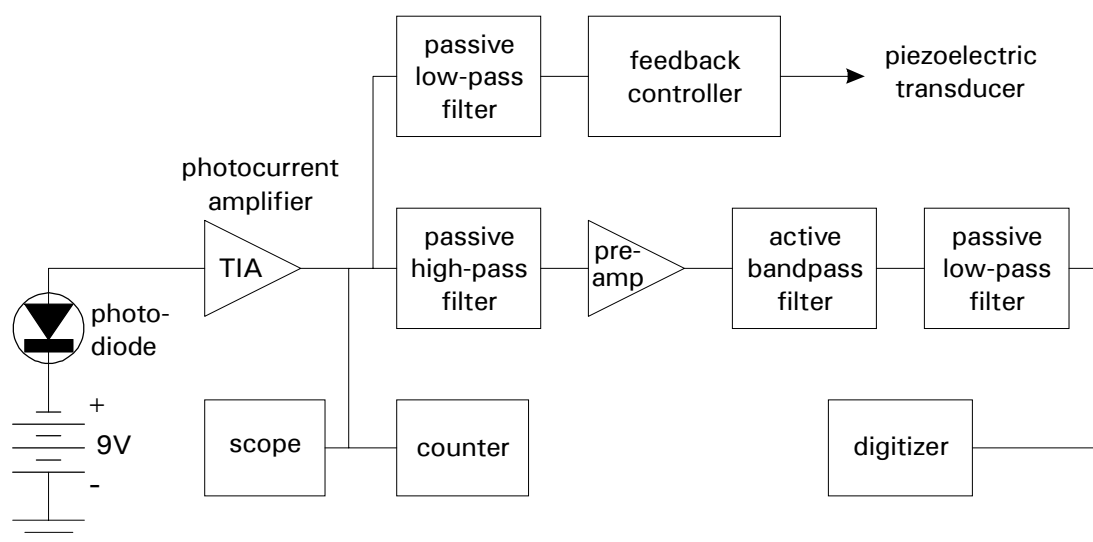


Figure 3.12. Signal acquisition and conditioning. See text of Section 3.4 for details.

3.4 Signal Conditioning and Acquisition

The output of the photocurrent (transimpedance) amplifier is the sum of an audiofrequency voltage due to vibrations of the oscillator and a near-dc level that depends on slow variations of the distance of the fiber face to the silicon oscillator. As shown in Figure 3.12, this output is split. One arm is low-pass filtered, and this near-dc component (below 0.1 Hz) is sent to the input of a feedback circuit, which is designed to maintain a setpoint voltage by driving a piezoelectric stack connected to the fiber support. In the other output arm, the audio-frequency component is high-pass filtered to remove the dc offset and to attenuate a large ~ 30 Hz interference that arises from acoustic noise exciting a spurious mechanical resonance in the apparatus. The resulting signal is amplified by a factor of 50 or 100 by a preamplifier (Stanford Research, Inc. Model SR552) and then filtered by a

rack-mount bandpass filter (Wavetek Model 442 Dual HI/LO filter). This filter is convenient because its passband can be set with front-panel switches. However, since its output includes some high-frequency electronic noise, its output is filtered through a two-pole low-pass passive (*LC*) filter before digitization. Digitization takes place inside a Pentium-based computer at 2 or 4 kilosamples per second and 16 bits vertical resolution using an ISA-compatible digitization board (Computer Boards, Inc. model CIO-DAS1602). Typically a one- or two-second transient is recorded when the oscillator is driven resonantly by forces exerted by the sample's modulated magnetization.

3.5 Rf System

Figure 3.13 shows the system of synthesizers, mixers, and amplifier that are used to deliver pulses of rf magnetic fields to the sample. The radio-frequency source for the experiments is a Signatec AWG502 arbitrary waveform generator board inside a Pentium-based IBM-PC-compatible computer. The board "clocks out" preprogrammed voltages on two channels at up to 50 megasamples per second, and it has a 64-kbyte data memory and a 256-byte program memory that allows some flexibility in looping of pulse programs. All the pulse sequences described in this thesis are computed and synthesized in real time by a C-language control program between iterations of rf application. These pulse programs are sent to the AWG board where they await a software trigger signal. The output of the AWG board is 8 V peak-to-peak at maximum in the frequency range 2-5 MHz (limited by the Nyquist frequency of the 10 MSample/s rate used in most of the experiments).

This 3 MHz range covers both the proton and fluorine Larmor frequencies in the 0.66 T field when mixed up to the ~ 28 MHz region with a local oscillator (Programmed Test Sources, Inc. PTS500).

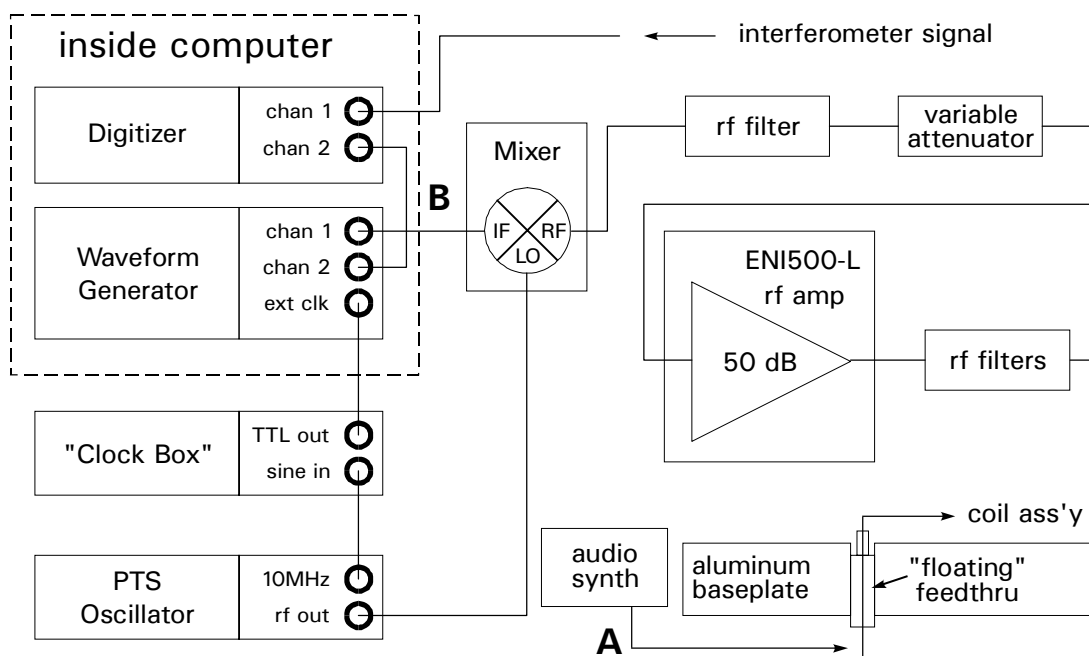


Figure 3.13. RF system as described in the text of section 3.5. At **A**, the NMR coil, which receives filtered rf current from the amplifier during experiments, can also be disconnected from the rf and connected to an audio-frequency synthesizer. This allows application of current at the sensor oscillator's resonance frequency, which can be used to calibrate the interferometric measurement of the oscillator's displacements. At **B**, a sinusoidal reference signal is sent from the AWG to a second channel on the digitizer. This timing channel acts as a phase reference for the detected interferometer signal.

The mixer unit is a collection of Mini-Circuits, Inc., devices configured for single-sideband operation by Lou Madsen². Its output is filtered and attenuated, followed by amplification by 50dB (ENI 3100L or 5100L-NMR). The output from the amplifier is again filtered and then transmitted through a floating-ground feedthrough to the interior of the vacuum bell jar, where it is transmitted through

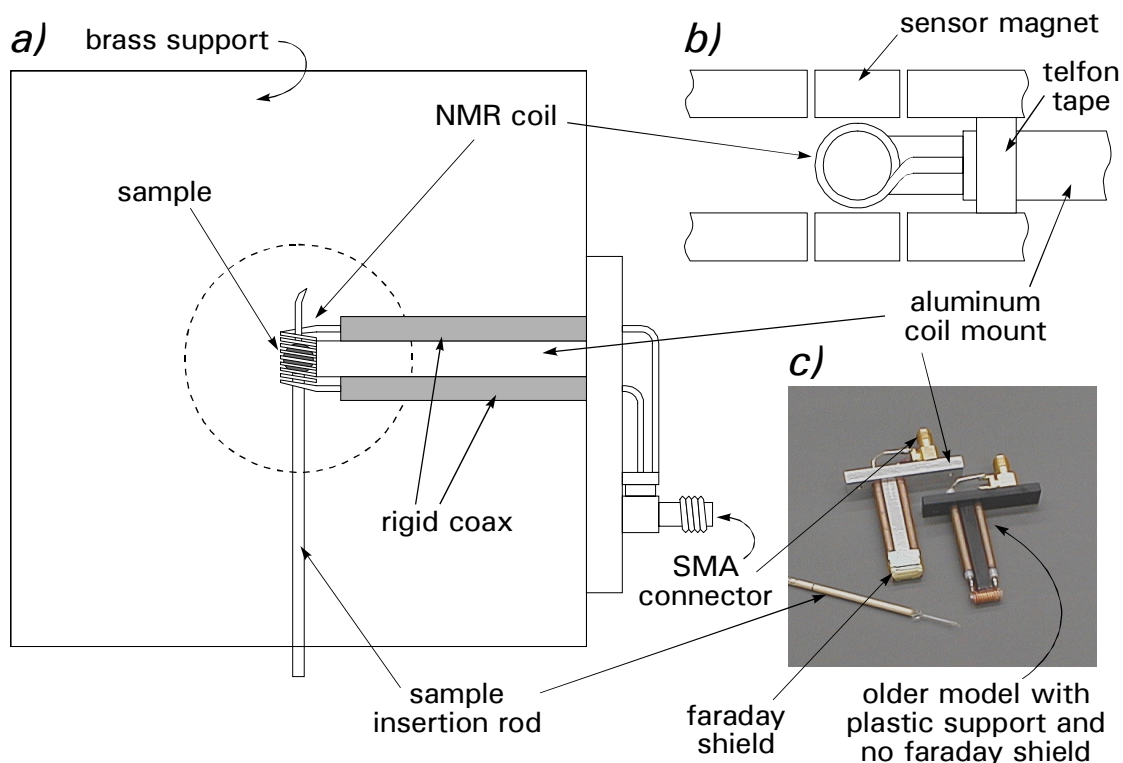


Figure 3.14. Therefore coil assembly. a) Top view. The sample is slid into the coil after the coil mount is installed as shown. b) Side view close-up showing coil underneath sensor magnet. c) Photo of two coil mounts and a liquid sample glued to an insertion rod.

~12 cm of micro-coax cable to an SMA connector on the NMR coil assembly shown in Figure 3.14.

Both the live and rf-ground voltages are brought into the magnet array through shielded rigid coax. The rigid coax that is used has outer and inner conductors made of a non-magnetic, low-permeability copper to avoid distortion of the static magnetic field. The ends of the rigid coax connect to a 9-turn coil that is wound within a beryllia tube and protected by an electric-arc suppressing mixture (GC Electronics Red GLPT Insulating Varnish, diluted with a few drops of toluene for smoother flow around small conductors). In order to get the sample as close as

possible to the moving sensor magnet, the wires of this coil (which are initially 32 gauge) are flattened to 150 μm thick between metal rollers. The beryllia tube has a high thermal conductivity, and in combination with the aluminum support to which it is bonded, it serves to conduct away heat that can change the resonance frequency of the oscillator. The outside of the beryllia tube is covered with a 25 μm thick, grounded Faraday cage to suppress the direct electromagnetic interaction of the coil with the sensor oscillator.

After the magnet assembly is put together², the coil assembly is inserted into the magnets so that the coil is directly under the moving sensor magnet. Then the sample, which is typically a liquid in a spherical bulb or a powder packed into a cylindrical tube between magnetically inert plugs, is affixed to a wooden dowel and inserted from another side, along the coil axis (see Figure 3.14).

The NMR coil is usually untuned. Our rf amplifier is strong enough to supply the necessary rf power into an untuned coil at the prototype size scale. This will also easily be the case for microfabricated BOOMERANG devices for low-power remote spectroscopy applications^{2,8}. The lack of tuning capacitors, in combination with a direct-digital approach to rf synthesis, allows us to apply multiple pulses on two different nuclei simultaneously on the same channel and also to apply the broad-band frequency-swept pulses (cyclic adiabatic passage) that drive the oscillator during detection with minimal amplitude modulation of the rf current in the coil, which is important to suppress spurious driving of the oscillator.

3.6 *Cyclic Inversion with Phase-Cycled Efficient ARP*

We have repeatedly referred to cyclic inversion of the sample's magnetization as the means by which oscillating forces are applied to the sensor oscillator in BOOMERANG. This inversion is so important to BOOMERANG that we shall consider it now in some detail. In order to drive the oscillator into mechanical resonance, the nuclear magnetization of the sample is inverted twice per oscillator period τ . This could be done with a train of π -pulses. However, even though the field is designed to be homogeneous, there is still a residual spread of Larmor frequencies, which as we have seen is of order 20 kHz. In order to optimize signal power, the oscillator driving procedure must repeatedly invert these inhomogeneously broadened spins with negligible loss in magnetization over and above losses due to unavoidable relaxation in a time period $\sim T_{1a}$.

We may quantify the needed efficiency of a single inversion process by observing that if each pass loses a fraction ε of the magnetization (in addition to the fraction $\tau/2T_{1a}$ lost to relaxation), then the exponential relaxation time will be decreased from T_{1a} to $\left(\frac{2\varepsilon}{\tau} + \frac{1}{T_{1a}}\right)^{-1}$. If this new relaxation time is to be more than half of T_{1a} , then ε must be less than $\tau/2T_{1a}$, which is 0.001 for $T_{1a} \approx T_1 = 1\text{ s}$ and $\tau \approx 1/500\text{ Hz} = 2\text{ ms}$. Thus we require that each pass be at least 99.9% efficient in inverting the entire sample magnetization. This limits use of the "sudden" approach of π -pulses (even composite pulses), which have limited spectral range given the required efficiency.

A better approach is adiabatic passage, which is known to be a more efficient means of inverting a population of inhomogeneously broadened spins. In adiabatic passage, the frequency $\omega(t) = \dot{\phi}(t)$ of the rf component of an applied field

$$\mathbf{B}(t) = B_0 \hat{\mathbf{z}} + B_1 (\hat{\mathbf{x}} \cos \phi(t) - \hat{\mathbf{y}} \sin \phi(t)) \quad (3.2)$$

is swept through Larmor resonance in a time that is short compared to spin relaxation but long compared to a period of spin precession in the “effective field” (see below) in a coordinate frame rotating at the frequency of the applied field. In Equation (3.2) and the following, we leave out a counter-rotating component that is present when linearly polarized rf is used. Let $\{\hat{\mathbf{i}}, \hat{\mathbf{j}}, \hat{\mathbf{k}}\}$ denote unit vectors in the rotating frame, with $\hat{\mathbf{k}}$ along the static field and $\hat{\mathbf{i}}$ along the instantaneous transverse rf field. These may be expressed in terms of laboratory-fixed unit vectors $\{\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}\}$ as

$$\begin{pmatrix} \hat{\mathbf{i}} \\ \hat{\mathbf{j}} \\ \hat{\mathbf{k}} \end{pmatrix} = \begin{pmatrix} \hat{\mathbf{x}} \cos \phi(t) - \hat{\mathbf{y}} \sin \phi(t) \\ \hat{\mathbf{y}} \cos \phi(t) + \hat{\mathbf{x}} \sin \phi(t) \\ \hat{\mathbf{z}} \end{pmatrix}. \quad (3.3)$$

The magnetization $\mathbf{M}(t)$ evolves according to the torque equation $\dot{\mathbf{M}} = \gamma \mathbf{M} \times \mathbf{B}(t)$, which reads

$$\dot{\mathbf{M}} = \gamma \mathbf{M} \times \left(B_1 \hat{\mathbf{i}} + \frac{\omega(t) - \omega_0}{\gamma} \hat{\mathbf{k}} \right) \quad (3.4)$$

in the rotating frame, with γ denoting the magnetogyric ratio of the spins. The equation of motion takes the form of Equation (3.4) even when the offset

$\Delta = \omega - \omega_0$ from the Larmor frequency $\omega_0 = \gamma B_0$ is time-varying⁹. The effective field in the rotating frame,

$$\mathbf{B}_{\text{eff}} \equiv \left(B_1 \hat{\mathbf{i}} + \frac{\omega(t) - \omega_0}{\gamma} \hat{\mathbf{k}} \right) = \left(B_1 \hat{\mathbf{i}} + \frac{\Delta(t)}{\gamma} \hat{\mathbf{k}} \right), \quad (3.5)$$

is a vector that traverses a path in the i - k plane as the instantaneous frequency of the applied field is changed. This path is a line parallel to $\hat{\mathbf{k}}$ if the magnitude of the transverse field is held fixed, as shown in Figure 3.15. If the angle $\theta(t)$ that the effective field makes with the transverse plane changes slowly compared to the precession rate:




Figure 3.15. Efficient ARP. In the rotating frame, the offset frequency $\Delta(t)$ is swept so that the angle $\theta(t) = \tan^{-1}(\Delta(t)/\gamma B_1)$ changes linearly in time.

$$|d\theta/dt| \ll \sqrt{\gamma^2 B_1^2 + \Delta^2}, \quad (3.6)$$

then the approximate solution to Equation (3.4) is a magnetization vector $\mathbf{M}(t)$ that follows the effective field vector $\mathbf{B}_{\text{eff}}(t)$. This is the basis of the so-called adiabatic rapid passage (ARP): the transverse field is applied at a frequency which begins far off resonance and is swept through the resonance frequency to a point far on the other side, inverting the magnetization.

Simple trigonometric considerations evident in Figure 3.15 make the rate of change in θ for a given change in Δ maximum when $\Delta=0$. This is also the condition

under which the adiabaticity condition (3.6) is most stringent. Therefore, both linear ($\Delta(t) \propto t$) and sinusoidal ($\Delta(t) \propto \sin \omega_m t$) sweeps in frequency place unnecessarily demanding constraints on the rate at which the magnetization can be efficiently swept through resonance, since in both these sweeps $|d\theta/dt|$ is largest just when \mathbf{B}_{eff} is weakest.

This was recognized by Hardy *et al.*^{10,11}, who pointed out that since the precession rate is very large when the offset frequency Δ is large, the sweep rate in ARP can be much faster at the beginning and end of the sweep, that is, when the rf is far from resonance. They demonstrated a far more efficient, tangent-based adiabatic rapid passage using a frequency sweep of the form

$$\omega(t) = \omega_0 - \omega_s \tan(\omega_s \alpha t), \quad (3.7)$$

where ω_s is a sweep shape parameter near $\omega_1 = \gamma B_1$ and where

$$\alpha = \left(\frac{2}{\omega_s T_0} \right) \tan^{-1} \left(\frac{\Omega}{\omega_s} \right). \quad (3.8)$$

T_0 is the total sweep time ($-T_0/2 \leq t \leq +T_0/2$), and 2Ω is the range of the angular frequency sweep. These pulses use far less rf power for given ranges of both offset frequency Δ and rf field inhomogeneity ΔB_1 for a given required inversion efficiency.

Such frequency-modulated pulses are created by direct digital synthesis straightforwardly. To do this, one calculates the phase $\phi(t)$ as a function of a discrete time variable and then stores values proportional to $\sin \phi(t)$, which are “clocked out” to the waveform generator when needed. The required phase is the time integral of the frequency in Equation (3.7):

$$\phi(t) = \int_{-T_0/2}^t \omega(t) dt = \omega_0(t + T_0/2) + \frac{1}{\alpha} \ln \frac{\cos \omega_s \alpha t}{\cos \omega_s \alpha T_0/2}. \quad (3.9)$$

The inversions could be done by sweeping back-and-forth through the NMR line. However, if there is any difference between back-and-forth sweeps in the non-NMR response of the oscillator to the rf power in the coil (such as heating, mechanical expansion of the coil, *etc.*), then the predominant Fourier component of the resulting force will be exactly at the oscillator frequency. This back-and-forth




Figure 3.16. Driving the oscillator with efficient ARP. The frequency of applied rf is swept according to equation 3.7, and there are two sweeps per oscillator period τ . Both sweeps begin on the high-frequency side of the center frequency ω_0 . The rate $d\omega_{rf}/dt$ is largest when the magnetization \mathbf{M} is nearly along the static field direction. The oscillator amplitude $z(t)$ is 90° out of phase with the driving force, which is proportional to M_z .

sweeping protocol can thus lead to spurious driving of the oscillator. We suppress these effects by sweeping from the same side of the NMR line on every inversion as shown in Figure 3.16. This ensures that such non-NMR forces are at multiples of *twice* the oscillator frequency instead of exactly on mechanical resonance.

A possible complication with such same-side sweeping is that for finite sweep width, an individual inversion ends with the effective field (and therefore the magnetization) having a non-vanishing transverse component. In order to prevent loss of this magnetization between passes, a 180-degree phase shift is applied on every other sweep. Figure 3.17 details how this phase shift corrects for such

imperfect inversion and retains more of the magnetization. We have found empirically that such phase cycling can prolong the driving interval by as much as a factor of four².

Figure 3.18 shows the result of inverting proton magnetization in a 2.6 mm diameter liquid water sample over a 1-second interval with the phase-cycled tangent sweeps. By integrating the harmonic-oscillator equation of motion

with an exponentially damped sinusoidal driving force, we can find the expected trajectory of the oscillator. When the oscillator's damping time τ is significantly shorter than the relaxation time T_{1a} of the magnetization, the result is

$$z(t) = Ae^{-t/T_{1a}}(1 - e^{-t/\tau})\cos \omega t, \quad (3.10)$$

where ω is the oscillator's resonance frequency, which is also the driving frequency (504 Hz for the transient in Figure 3.18). The amplitude A , which can be found using the theory of Chapter 2, is equal to the on-resonance steady-state amplitude were the magnetization not to decay.




Figure 3.17. Phase cycling for efficient ARP. After one pass through resonance, the magnetization \mathbf{M} is nearly, but not exactly, inverted (1). If the effective field at the start of the next pass through resonance is applied at (2) in this rotating-frame picture, then \mathbf{M} is not perfectly aligned with the effective field. A component of \mathbf{M} transverse to the effective field will be dephased. The net effect on \mathbf{M} is a fractional loss of $\varepsilon = 1 - \cos 2\theta \approx 2\theta^2$. A 180-degree phase shift of the rf, which places the effective field at (3), mitigates this loss.

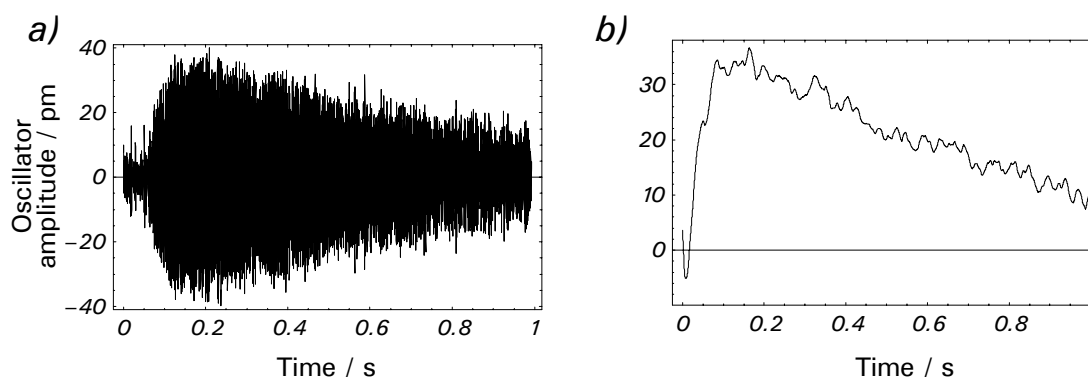


Figure 3.18. Driving the oscillator. a) Time series recorded during cyclic inversion of proton magnetization in a 2.6-mm diameter liquid water sample. b) Envelope of the signal in figure (a). Data from (a) were digitally mixed down to DC and then filtered in order to fit the data. The transient excitation was fit to equation (3.10) in order to find T_{1a} (0.687 s), τ (51 ms), and the amplitude (44.7 pm). The measured amplitude agrees with predictions from the theory in Chapter 2 (45.1 pm) to within about 1%.

Figure 3.18 b shows the amplitude envelope of the transient in Figure 3.18 a. A fit to Equation (3.10) yields $\tau=51$ ms, $T_{1a}=0.687$ s, and $A=44.7$ pm. Agreement of the amplitude with expectations (45.1 pm) using an estimate of the sensor's magnetization from the observed Larmor frequency and the signal theory of Chapter 2 is remarkable ($<1\%$). The measured T_{1a} corresponds to about 693 inversions by the time the magnetization decays by the factor e . This relaxation time is substantially shorter than the T_1 of water² (4.3 s) as measured by inversion-recovery (see Chapter 4). These results indicate that the whole sample magnetization is being used to drive the oscillator, but that substantial relaxation of the magnetization is taking place as a result of the inversions.

3.7 System Noise Analysis

3.7.0 Brownian motion revisited. Displacement calibration

Even at the prototype size scale, the principal source of noise in the recorded transients is due to the Brownian motion of the oscillator. This can be quantified in two ways – as noise in the driving force or noise in a displacement measurement. On general thermodynamic grounds^{7,12}, the rms fluctuation in the average force on the oscillator is predicted to be

$$F_{z,rms} = \sqrt{4k_B T \alpha \Delta f} . \quad (3.10)$$

This may be viewed as the square root of a force-noise spectral density

$$S_F^{1/2} = \sqrt{4k_B T \alpha} \quad (3.11)$$

multiplied by the square root of a bandwidth, Δf . If the bandwidth of the measurement is substantially less than the bandwidth $\gamma/2\pi = 1/4\tau$ of the mechanical oscillator (which is true for the pointwise detection schemes of Chapter 4 when $T_{1a} \gg \tau$), then the oscillator is approximately in steady state during the time the oscillator is driven. In this case the square root of the corresponding displacement-noise spectral density at mechanical resonance is obtained from (3.11) by multiplying $S_F^{1/2}$ by the quality factor $Q = \omega/\gamma$ and dividing by the spring constant $m\omega^2$:

$$S_x^{1/2} = \frac{Q}{m\omega^2} \sqrt{4k_B T \alpha} = \frac{1}{m\omega\gamma} \sqrt{4k_B T \alpha} = \sqrt{4k_B T / m\omega^2 \gamma} . \quad (3.12)$$

With $m = 139 \text{ mg}$, $\omega/2\pi = 496 \text{ Hz}$, and $\gamma/2\pi = 5.0 \text{ Hz}$, $S_F^{1/2}$ and $S_x^{1/2}$ are $8.4 \text{ pN}/\sqrt{\text{Hz}}$ and $620 \text{ fm}/\sqrt{\text{Hz}}$ ($6.2 \text{ mÅ}/\sqrt{\text{Hz}}$), respectively, at room temperature (293 K). As we shall see below, this contribution dominates other noise sources, but only by a small factor at the prototype size scale, and so some care was exercised in suppressing these other sources as much as possible. In smaller spectrometers, the Brownian-motion noise will be fractionally larger relative to these other sources of instrument noise. For direct comparison to Brownian-motion noise, each noise source is referred back to a displacement noise spectral density through multiplication by relevant gain factors, whose nominal values are shown in the signal path of Figure 3.19.

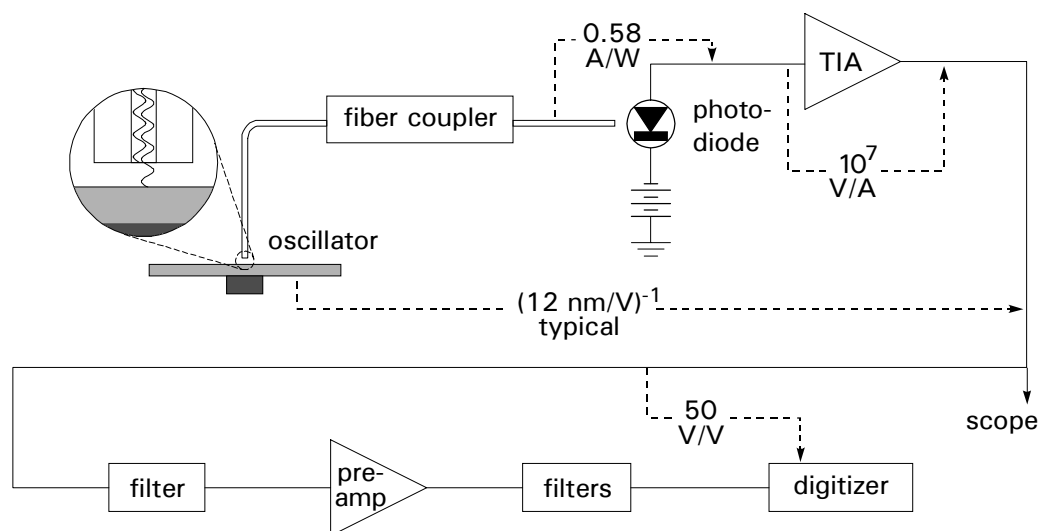


Figure 3.19. Signal conditioning path. Selected gain factors are shown for important connections in the signal path.

3.7.1 Photon shot noise

When the interferometer is set to its most linear operating point, there is a dc component to the light incident on the detector. The dc level of the light intensity at the photodetector is the source of two kinds of noise. The first is shot noise, which is due to the Poisson statistics of independent “photon arrivals,” each with energy $h\nu$, from the light field¹³. If P denotes the optical power incident on the detector, then in a time interval Δt , $N = P\Delta t/h\nu$ photons arrive at the detector on average. The rms fluctuation in this average is \sqrt{N} , and therefore the rms fluctuation in the optical power is

$$P_{\text{rms}} = \frac{h\nu}{\Delta t} \sqrt{N} = \sqrt{Ph\nu/\Delta t} = \sqrt{2Ph\nu\Delta f}, \quad (3.13)$$

where explicit dependence on the the time interval Δt has been suppressed in favor of the corresponding bandwidth $\Delta f = 1/2\Delta t$. This allows convenient comparison of the square root of the corresponding noise spectral density,

$$S_{h\nu}^{1/2} = \sqrt{2Ph\nu}, \quad (3.14)$$

to other noise sources, including Brownian motion. With $P = 1.7 \mu\text{W}$ and $\nu = c/780 \text{ nm}$, $S_{h\nu}^{1/2} = 930 \text{ fW}/\sqrt{\text{Hz}}$. We convert this value to the square root of the corresponding displacement-noise spectral density by multiplying (3.14) by the detector responsivity (0.58 A/W) to obtain the corresponding photocurrent noise, then by the transimpedance ($10^7 \Omega$) to obtain the corresponding voltage noise out of the photocurrent amplifier, and finally by the displacement sensitivity (typically 12 nm/V). The result is $S_{x,h\nu}^{1/2} = 65 \text{ fm}/\sqrt{\text{Hz}}$.




Figure 3.20. Interferometer displacement calibration. a) Voltage at the output of the transimpedance amplifier observed on an oscilloscope registers the voltage span associated with the interferometer's trough-to-crest displacement change. b) The laser intensity goes from trough to crest when the fiber-to-oscillator gap increases by $\lambda/4 = 195 \text{ nm}$ because the wave must traverse the gap twice before it re-enters the fiber and mixes with the reference wave. The displacement sensitivity may thus be calculated using equation (3.16).

The displacement sensitivity, the proportionality constant relating displacement to observed voltage at the output of the photocurrent amplifier, is used here because it is readily measured in the following way. The oscillator is driven to an amplitude that is large enough to observe the “folding over” of the voltage-to-distance relation (see Figure 3.20). These nanometer-scale oscillations are well within the expected range of linearity of the silicon oscillator, and so the nonlinearity of the observed voltage is due entirely to “interferometer action.” The difference between the extrema of the oscilloscope trace corresponds to the voltage difference associated with a displacement of $\lambda/4$. Since the slope of the voltage-to-displacement sine curve is greater than the ratio of voltage span to displacement by $\pi/2$ in the linear region (as shown in Figure 3.20), the displacement sensitivity is

$$\frac{2}{\pi} \times \frac{\lambda/4}{V_{\text{max}} - V_{\text{min}}} \quad (3.16)$$

3.7.2 Photon pressure fluctuations

The second type of noise due to the Poisson statistics of the laser radiation is the result of fluctuations in radiation pressure. These can be estimated with a simple model that assumes that a single photon impact on the oscillator transfers $2h\nu/c$ to the oscillator's momentum. Again, if in a time Δt there are $N = P\Delta t/h\nu$ photon impacts, then the rms fluctuation in the transferred momentum is

$$p_{\text{rms}} = \frac{2h\nu}{c} \sqrt{N} = \sqrt{4Ph\nu\Delta t/c^2} . \quad (3.17)$$

If this momentum fluctuation takes place in a time Δt , then it may be viewed as a random force $p_{\text{rms}}/\Delta t$, which has the force spectral density

$$S_{F,hk}^{1/2} = \sqrt{8Ph\nu/c^2} \quad (3.18)$$

and, in analogy with Equation (3.12), the displacement-noise spectral density

$$S_{x,hk}^{1/2} = \frac{1}{m\omega\gamma} \sqrt{8Ph\nu/c^2} . \quad (3.19)$$

At the BOOMERANG prototype size scale, this contribution ($5.4 \times 10^{-7} \text{ fm}/\sqrt{\text{Hz}}$) is totally negligible. While it will become more important in force-detected NMR at reduced size scales¹⁴, it is not a limiting factor in any proposed designs.

3.7.3 Photocurrent shot noise

Photocurrent shot noise is due to the Poisson statistics of discrete charge carriers in the electronic current through a diode. As this noise source is directly

correlated with the statistics of photon arrivals, we have already accounted for it in analyzing the photon shot noise.

3.7.4 Johnson noise in the transimpedance

The photocurrent amplifier uses a resistor at the input to convert the current to a voltage. The Johnson noise in this resistance forms the bulk of the noise added by this amplifier. The noise specification in the amplifier's documentation¹⁵ is consistent with this fundamental physical argument, and the amplifier was found to behave according to specification. The square root of the voltage-noise spectral density due to Johnson noise in a resistance R is given by

$$S_{\Omega}^{1/2} = \sqrt{4k_{\text{B}}TR}. \quad (3.20)$$

For a transimpedance of $10^7 \Omega$, this is $S_{\Omega}^{1/2} = 400 \text{ nV}/\sqrt{\text{Hz}}$. This can be converted to the square root of a displacement-noise spectral density by multiplication by the displacement sensitivity. The result is $S_{x,\Omega}^{1/2} = 4.8 \text{ fm}/\sqrt{\text{Hz}}$.

3.7.5 Electronic noise in the preamp

The Stanford Research preamp documentation¹⁶ specifies the voltage noise of the preamp referred to the input as

$$S_{\text{E}}^{1/2} = 1.5 \text{ nV}/\sqrt{\text{Hz}}. \quad (3.21)$$

Multiplying by the displacement sensitivity, we find $S_{x,\text{E}}^{1/2} = 0.018 \text{ fm}/\sqrt{\text{Hz}}$, which is negligible compared to Brownian motion.

3.7.6 Digitization noise

When the analog signal is digitized, the discretization of the signal introduces noise into each sample of the recorded data. This is shown in Figure 3.21. The added noise in a given voltage sample will be distributed uniformly over the range $+\epsilon/2$ to $-\epsilon/2$, where ϵ is the step size of the digitization process. This corresponds to an rms average fluctuation of $\epsilon/\sqrt{12}$ in the bandwidth $\Delta f = 1/2\Delta t$, where Δt is the sampling time. The relevant voltage-noise spectral density (at the digitizer) is




Figure 3.21. Digitization noise. In a given time slice, the requirement that the output of the digitizer be only one of a set of fixed values adds uncertainty to the measured value. This uncertainty is centered at zero and uniformly distributed over a range equal to the step size of the digitization process.

$$S_D^{1/2} = \frac{\epsilon}{\sqrt{12}} \sqrt{2\Delta t} = \epsilon \sqrt{\Delta t/6}. \quad (3.22)$$

For the prototype, a 16-bit digitizer was used, with full range -1.25V to $+1.25\text{V}$, corresponding to $\epsilon = 38.1\mu\text{V}$. For samples acquired every $500\mu\text{s}$, $S_D^{1/2} = 348\text{nV}/\sqrt{\text{Hz}}$ at the digitizer. To convert this to a corresponding displacement noise spectral density, this must be divided by the gain of the preamp (typically 50) and multiplied by the displacement sensitivity, yielding $S_{x,D}^{1/2} = 0.084\text{ fm}/\sqrt{\text{Hz}}$. This estimate is for an otherwise noiseless digitizer. In the prototype experiments, it was observed that even with the inputs of the digitizer grounded, a frequency-independent electrical noise corresponding to ~ 3 bits peak-to-peak (rms amplitude

$2^3 \epsilon / 2\sqrt{2}$) is recorded in the data, probably caused by electrical noise inside the computer. Even with this factor, the digitizer added negligible noise compared to Brownian motion.

References

- 1 G. W. Castellan, *Physical Chemistry*, 3rd ed. (Addison-Wesley, Reading, MA, 1983).
- 2 L. A. Madsen, Ph. D. Thesis, California Institute of Technology, 2002.
- 3 Magnetic Alloys Catalog (Carpenter Technology Corp., 1995).
- 4 G. M. Leskowitz, L. A. Madsen, and D. P. Weitekamp, *Sol. St. Nucl. Magn. Reson.* **11**, 73 (1998).
- 5 W. D. Pilkey, *Formulas for Stress, Strain, and Structural Matrices* (John Wiley & Sons, Inc., New York, 1994).
- 6 D. Rugar, H. J. Mamin, and P. Guethner, *Appl. Phys. Lett.* **55**, 2588 (1989).
- 7 C. W. McCombie, *Rep. Prog. Phys.* **16**, 266 (1953).
- 8 T. George, et al., in *IEEE Aerospace Conference* (Big Sky, Montana, 2001).
- 9 C. P. Slichter, *Principles of Magnetic Resonance*, 3rd ed. (Springer-Verlag, Berlin, 1990).
- 10 C. J. Hardy, W. A. Edelstein, and D. Vatis, *J. Magn. Reson.* **66**, 470 (1986).
- 11 C. J. Hardy and W. A. Edelstein, *J. Magn. Reson.* **69**, 196 (1986).
- 12 S. Chandrasekhar, *Rev. Mod. Phys.* **15**, 1 (1943).
- 13 C. W. Gardiner, *Handbook of Stochastic Methods for Physics, Chemistry, and the Natural Sciences*, 2nd ed. (Springer, Berlin, 1985).
- 14 J. A. Sidles, J. L. Garbini, K. J. Bruland, D. Rugar, O. Züger, S. Hoen, and C. S. Yannoni, *Rev. Mod. Phys.* **67**, 249 (1995).
- 15 Model 181 Current Sensitive Preamplifier manual (Princeton Applied Research Corp., 1978).
- 16 Model SR530 Lock-In Amplifier manual (Stanford Research Systems, 1989).