TEACH SPIN
Instruments Designed for Teaching

Noise Fundamentals

NF1-A

A PRODUCT OF TEACHSPIN, INC.

Written by David Van Baak
Edited by George Herold
Jonathan Reichert

Adapted for ATI by
Johannes Majer, Stefan Putz v.1.1
Christian Koller v.1.2

TeachSpin, Inc.
2495 Main Street   Suite 409   Buffalo, NY   14214-2153
Phone: (716) 885-4701
Fax: (716) 836-1077
www.TeachSpin.com

0. Introduction

0.0. Parts of this apparatus

Figure 0.0 shows the major components of the experiment. If anything displayed here is missing or not functional, please contact your advisor immediately.

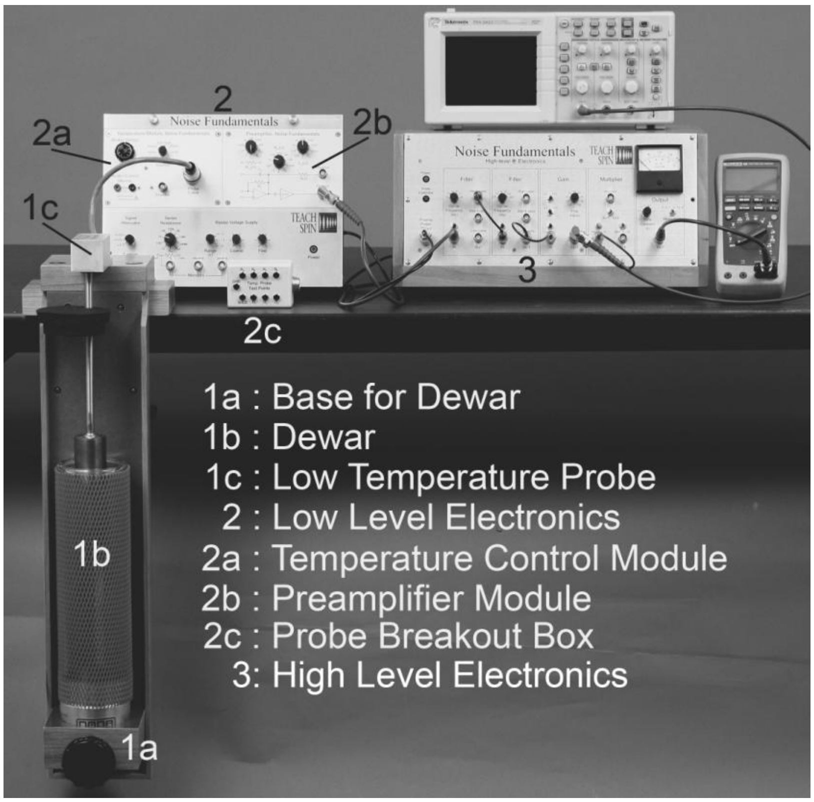

Figure 0.0: The parts of Noise Fundamentals apparatus

The main two parts are the high level electronics in a wooden case (3) and the low level electronics in a steel and aluminum case. As you see there is also a low temperature setup, which you will need for the final experiment of this practicum. This setup includes a variable-temperature sample probe, a Dewar vessel, the Dewar support and a breakout box. In addition there are two plastic parts boxes containing various tools and spare parts (not shown), which you might need for the optional part.

The measurements are mostly performed by using the provided digital voltmeter and the oscilloscope, which this experiment shares with the NMR setup. (You might have to take it from there, please be careful. In case there are any uncertainties, please contact your advisor.)

0.2. Definition, kinds, uses of noise

Just as a weed is an unwanted plant, a noise, *ordinarily speaking*, is an unwanted sound. In the fields of physics, electrical engineering, and many other places, we extend the definition of 'noise' beyond acoustics to the general field of information. Since almost any signal that's a function of time can be translated into a voltage, we will often use the concept of a voltage signal. We'll call it a 'noisy signal' if, in addition to the voltage we expect or wish to see, there is unwanted, typically (but not always) a randomly-fluctuating, voltage. Surprisingly, the noise signal is sometimes not only wanted, but is the essence of the measurement.

There are several kinds of noise. One of them is 'interference', which is the presence of an unwanted signal, added to the desired signal. It's easy to imagine that your neighbor's electronic apparatus is polluting your TV or radio signal with some sort of interference. The kind of interference students are likely to encounter in these experiments probably comes from three sources: electrostatic coupling to the apparatus from fluorescent lights in the laboratory, electromagnetic coupling due to nearby transformers or motors, and vibrational coupling due to microphonic components within the unit.

Another source of noise we will call 'technical noise' since it is the noise generated by the technique of the investigation, or that gets into the circuits due to *faulty* experimental techniques. For example, a failure to tighten the cover on the preamplifier section, or a poor electrical connection to the first-stage op-amp, can add extraneous noise to the signal path.

Of greatest interest to us is 'fundamental noise', noise that is intrinsic and inevitable because of the physical nature of an apparatus. We'll observe noise sources that arise from the Second Law of Thermodynamics, and from the quantization of electrical charge. Physicists and electrical engineers know these as Johnson and shot noise ("Schrotrauschen") respectively. Noise sources like this display the characteristics of non-periodic, unpredictable, random waveforms, but nevertheless conforming, in their statistical properties, to universal laws.

Fundamental noise is especially worthy of study, for at least two reasons. The first reason is that fundamental noise presents us with a physics-based limit on the degree to which we can measure in a given experiment. In many cases in research and technology, it often defines what is possible within the limits of physical law. In particular, fundamental noise can and does set limits to the rate of data-transfer in a host of contexts in communication.

The second reason we care about noise is that it becomes possible to use noise to measure the values of some fundamental constants. Boltzmann's constant k_B can be determined from the voltage or Johnson noise of resistors; and the magnitude of the charge on the electron, e, can be determined from the current or shot noise of a photocurrent.

But measurement of 'fundamental noise' has its experimental challenges. There is a saying about noise measurements: 'you're either measuring too much or too little signal'. You will understand this quip better after you have had some experience with these measurements. Our advice here is to read both the manual and some of the references and do your measurements carefully. **But most of all, have fun!**

0.3. Measurement Tasks & Protocol

During this practicum you will perform a number of measurements that should be recorded in a protocol. Please give comprehendible plots of the following measurements:

- Recording of Noise sample
- Variation of Noise (table)
- Dependence on Resistance (for all 3 internal resistors, external: optional)
- Dependence on bandwidth (choose twenty or more combinations)
- Calculation of Boltzmann's constant
- Interpolation of absolute zero

It is not necessary to give any derivation of the theory, but the plots should be interpreted and explained. In case additional questions are asked in the text, try to answer them. Give (where applicable) an estimate of the errors and where they come from.

1. Johnson noise at room temperature

1.0 The reasons for Johnson noise, and its predicted size

Everyone capable of a little electrical engineering and physics knows Ohm's law $V = i\,R$, which really says that there's a potential difference ΔV across any resistor R which has a current i passing through it. This of course predicts a ΔV of *zero* for a resistor with no current. But for deep reasons, any actual resistor at any temperature above absolute zero, will display a 'noise voltage' $V_J\,(t)$ across its terminals, a potential difference that has all the character of an internal (a.c.) *emf (electromotive force)* built into the resistor. The emf which the resistor generates is called 'Johnson noise', and it arises because of the deep thermodynamic connection between dissipation (which any resistor surely has) and fluctuations (which here show up as a fluctuating emf). The size of this emf is also predicted by fundamental theory, and it should not surprise you to learn that $V_J\,(t)$ is, *on average*, zero. But $V_J\,(t)$ exhibits fluctuations, positive and negative, about that average value of zero. To quantify these, we form the (always-positive) *square* of $V_J\,(t)$, and time-average that, giving a 'mean square' voltage which we denote as $< V_J^2\,(t) >$. The predicted value for $< V_J^2\,(t) >$ was first deduced by Nyquist, following Johnson's empirical discovery of the noise, and it's given by the expression

$$< V_J^2\,(t) > = 4\,k_B\,R\,T\,\Delta f.$$

Here k_B is Boltzmann's constant, T is the (absolute) temperature of the resistor, and Δf is the novel factor -- it is the 'bandwidth' used in the measurement electronics.

The involvement of bandwidth Δf is a first hint that 'noise' is quite distinct from 'signal'. Everyone starts with 'd.c. signals', which have nothing but a sign and a value, in Volts. Then there are 'a.c. signals', which have a magnitude (perhaps specified by amplitude, or rms value, or peak-to-peak excursion) but also a *frequency*, or a mixture of frequencies. But it is the essence of fundamental noise that it contains, or is composed of, *all frequencies*. In fact, the amount of energy we can get out of a 'noise source' depends on the *range* of frequencies to which we arrange to be sensitive, and this is the reason for the inclusion of the bandwidth-factor Δf in the expression above.

How large a Johnson-noise voltage should we expect from a typical resistor? Let's calculate this mean-square voltage for a 100-kΩ resistor at room temperature. Suppose that our electronics for detecting and measuring $V_J\,(t)$ are fully sensitive to all frequencies from 0 to 100 kHz, but entirely insensitive to higher frequencies. Then:

$$T = 22\,^{o}\mathrm{C} = 295\ \mathrm{K}$$
$$k_B = 1.38 \times 10^{-23}\ \mathrm{J/K}\ \text{(textbook value)}$$
$$\Delta f = 100\ \mathrm{kHz} = 10^5\ \mathrm{Hz}$$

$$\begin{aligned}<V_J^2(t)> &= 4\,(1.38 \times 10^{-23}\ \mathrm{J/K})\,(295\ \mathrm{K})\,(10^5\ \Omega)\,(10^5\ \mathrm{Hz}) \\ &= (1.63 \times 10^{-20}\ \mathrm{J})\,(10^5\ \mathrm{V/A})\,(10^5\ /\mathrm{s}) \\ &= 1.63 \times 10^{-10}\ \mathrm{V}^2.\end{aligned}$$

Not everyone is familiar with the curious unit of the square-of-a-Volt, so we often take the square root of this mean-square noise voltage, to give a 'root-mean-square' or 'rms' measure of the noise voltage,

$$V_J(\mathrm{rms}) \equiv \langle V_J^2(t)\rangle^{1/2} = 1.28 \times 10^{-5}\ \mathrm{V} = 12.8\ \mu\mathrm{V}.$$

So if we have a room-temperature 100-kΩ resistor simply hooked up to an ideal voltmeter, and if that voltmeter responds to all (but only) frequencies under 100 kHz, then the voltmeter's instantaneous reading will *not* be zero volts, but instead will fluctuate (rapidly: in this case, on a microsecond time scale) around zero, with typical excursions of order ±10 μV. We further assert that this is an actual emf intrinsic to the resistor, and it will still be present, though typically unwanted, in addition to any iR-drop that the resistor may exhibit. It follows that measurement of any iR-drop to microVolt precision in such a case would require thinking about this effect.

There are many textbook derivations of Nyquist's prediction, and the best of them emphasize the connection to thermodynamics and to blackbody radiation. Here's a 'thought experiment' to help you see that some sort of Johnson noise must exist. First imagine a cubic meter of iron at room temperature and another cubic meter of cold iron (say, at temperature T = 4 K), spaced 10 meters apart in empty space. (If you like, think of them as located at the two focal points of a large evacuated ellipsoidal reflecting cavity which surrounds them both, and isolates them from the external universe.) It should be clear to you that each iron block is giving off blackbody radiation, with a range of frequencies and in all directions -- but that the warm block is giving off lots more. Since the blackbody radiation of each block will run into the other block, there will be a net flow of (radiant) energy from the warmer block to the colder one, and their temperatures will therefore start to equilibrate.

Now imagine a 50-Ω resistor at room temperature, connected to nothing but a lossless coaxial cable of 50-Ω impedance; and imagine there's another 50-Ω resistor, but down in a Dewar at T = 4 K, connected to the far end of this cable. Even if there is no thermal conductivity in the cable, there is still electrical conductivity. It's the 'Johnson emf' in each resistor which still acts like a black-body source, here generating travelling waves of (confined) radiation along the one-dimensional cable structure, and that 'radiation' is caught and dissipated in the far end's resistor. This is the mechanism by which the two resistors will tend toward thermal equilibrium, as the hotter resistor will experience a net outflow, and the colder a net inflow, of electrical energy.

1.1 'Seeing' Johnson noise

This exercise will let you see, directly on an oscilloscope, a time-dependent waveform which can be traced all the way back to the Johnson noise generated in a resistor.

You need to plug, into your 100-to-240-V outlet, the line cord of the the universal power supply which supplies power to the high-level electronics (HLE). You should see a green LED on the transformer unit light up. Now connect the output of this supply to the receptacle on the back of the HLE. You should see a green LED on the front panel of the HLE light up. (Note there is no power switch in the HLE box; instead, it gets powered up as soon as you establish the power-supply connections.) Now find the power cable emerging from the LLE box, and plug it into the connector on the front panel of the HLE box. You should see a green LED light up on the front panel of the LLE. Once you have three green LEDs lit, everything in your system is being powered.

Set the switch to select a 'source resistor' of R_{in} = 100 kΩ in the pre-amplifier module installed in the LLE box. This resistor is connected only to the high-impedance input of the first stage of amplification in the pre-amp. That first stage is wired to give a 'gain', or amplification factor, of 6.00, *provided* you set the feedback resistor, R_{f}, to its 1-kΩ setting. (The feedback capacitance C_{f} is not connected in the default mode, so its setting is irrelevant.) Read the graphics on the panel of the pre-amp to see that there is an additional amplification stage, with gain 100, following this first stage. Now you can connect the pre-amp's output, by a coaxial cable, to an oscilloscope, to see if there is any signal present. Use a rather sensitive vertical scale on your 'scope (of perhaps 10 mV/division sensitivity), a sweep speed of 5 μs/div on the horizontal axis, and trigger near zero volts.

Below are the schematic, and the wiring, diagrams of the circuit you're using.

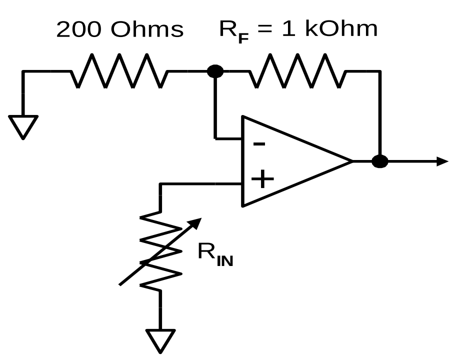

Fig 1.1a: Johnson noise preamplifier schematic

The wiring diagram for this configuration is shown in Figure 1.1b. The connections indicated in grey-scale printing are those you need to check, or establish. (By contrast, connections shown in thin solid lines are already established for you on the printed-circuit boards.)

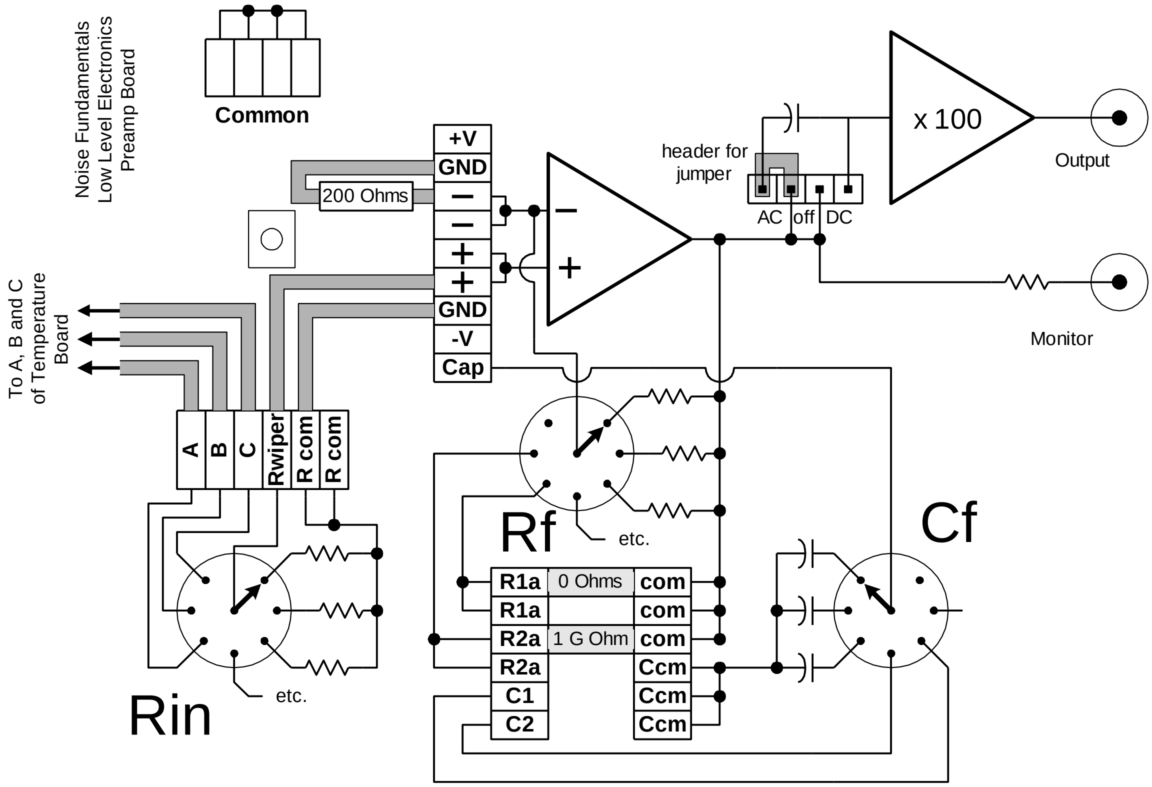

Fig. 1.1b: Wiring diagram of the default condition of the interior of the low-level electronics.

The signals you've seen emerging from the pre-amp are rather small. So next use a BNC cable to convey the pre-amp output to the HLE box instead, where you can filter and amplify the still-small noise signals. If you use the settings and the cabling shown in Fig.1.1c, you will be selecting a frequency band, extending from about 100 Hz to about 100 kHz, to pass along to the main amplification stages. The first filter shown has its high-pass output in use; you may think of this as passing frequencies on the high side of 100 Hz, or equivalently as blocking frequencies below 100 Hz. The second stage is used as a low-pass filter, here passing all frequencies on the low side a chosen 100 kHz. So after the output of the two filters, you have Johnson noise, pre-amplified by factor 600., and then filtered to pass only the 0.1 - 100 kHz frequency band.

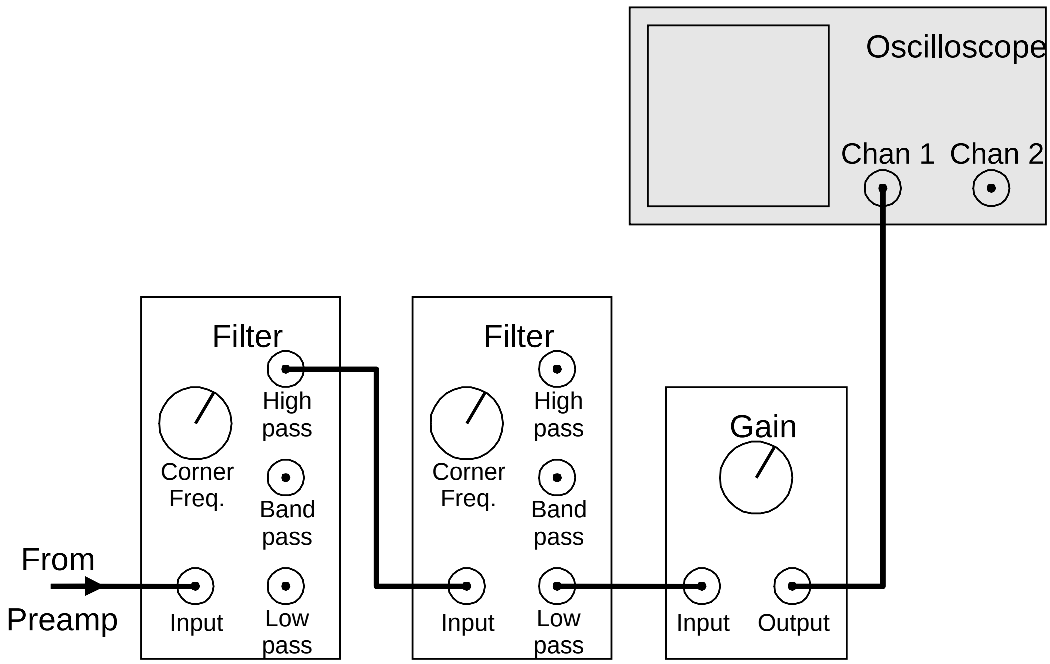

Fig. 1.1c: Cabling diagram for first use of the high-level electronics.
(left) Filter: selector to .1k (for 0.1 kHz), switch to AC (for a.c. coupling)
(right) Filter: selector to 100k (for 100 kHz), switch to AC (for a.c. coupling)
Gain Fine Adjust 30, toggle x1, toggle x10

Notice the figure shows more cabling, now to amplify this signal by a further factor of 300. You achieve this by a setting of gain x1 and x10 at two toggle-switch settings, and a further gain of x30 on the rotary switch setting. (Here too you can switch to AC for a.c. coupling at the input.) Finally, at the output of this main amplifier, you'll have a signal large enough to see easily on a 'scope. A view of it, using a 2 V/div vertical sensitivity, and a 10 μs/div horizontal scale, is shown in Figure 1.1d.

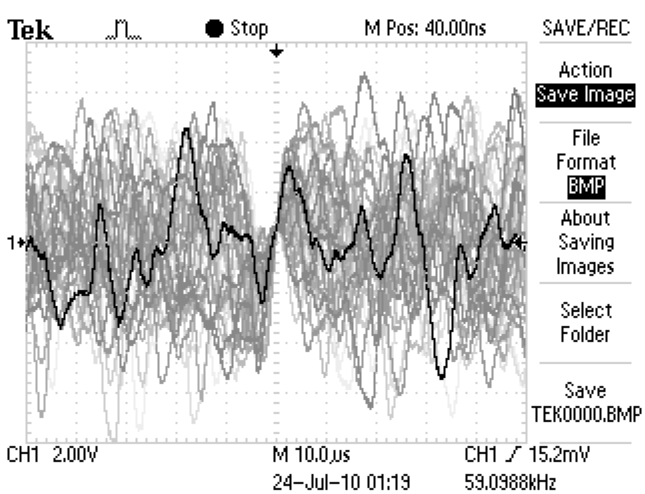

Fig. 1.1d: Samples of amplified Johnson noise from a 100-kΩ resistor, using pre-amp gain 600, filtering to 0.1 - 100 kHz bandwidth, and main-amp gain 300. Vertical scale 2 V/div, horizontal scale 10 μs/div, triggering on positive-going zero-crossings.

To get a first, qualitative, indication that this 'noise signal' has something to do with the original source resistor at the front end of this pre-amp/filter/main-amp chain go back to the pre-amp, and change the source resistor from 100 kΩ to 10 kΩ. You should see the size of the noise signal on your 'scope *change* -- it should decrease, and by a factor of about three.

For a first rough understanding of the *size* of these 'scope signals, consider our claimed 13-μV (rms measure, in the 0-100 kHz band) Johnson-noise signal emerging from a 100-kΩ source resistor. The pre-amp gain of 600 ought to raise this to about 8 mV (rms), and further main-amp gain of 300 ought to raise this to about 2.5 V (rms). (The intervening filter stages enforce the limitation to the 0.1 – 100 kHz band, and they provide a gain very near 1.00 within that band.) We'll see later a good way to measure the rms value of signals such as shown in Fig. 1.1d, but you can now see why those voltage excursions fall (mostly) in the ±5-V range.

If your signals differ dramatically from those shown here, something is amiss. (See Appendix A.6 for some suggestions about 'troubleshooting'.) It's certainly possible for the signal you see to be smaller, say if you've made wrong connections or wrong settings. It's also possible for the signal to be 'too large', particularly if there are unwanted (interference) signals present. (Appendix A.5 discusses interference, its possible sources, and cures.) But the apparatus you're using, in the configuration you've set up, ought to be displaying a noise almost wholly due to nothing else than the Johnson noise of your source resistors. It is the universality of Johnson noise that lets us be sure that your signals should match, in rms measure, those shown here, certainly to within a factor smaller than two!

1.2 Quantifying Johnson noise

If you've done section 1.1, you've seen a rapidly-fluctuating signal on an oscilloscope which we claim is due mostly to Johnson noise, and which you now want to quantify. The method we'll describe here executes quite directly, in analog electronics, the very operation built into the mean-square definition of noise. You need one more cable to convey the filtered-and-amplified noise signal to the Multiplier module, configured as a 'squarer' as shown in the Fig.1.2a. Conduct the noise signal to the 'A' input, and choose the AxA on the toggle switch. The multiplier circuit delivers at the MONITOR point, a real-time output voltage

$$V_{out}(t) = [V_{in}(t)]^2 \,/\, (10\text{ V})\ ,$$

which still has dimensions Volts (due to the fixed 'scale factor' of 10 Volts in the denominator above). Take a look at $V_{out}(t)$ on your 'scope, and notice that it is always positive, *un*like your input noise signal $V_{in}(t)$, which is as often negative as positive.

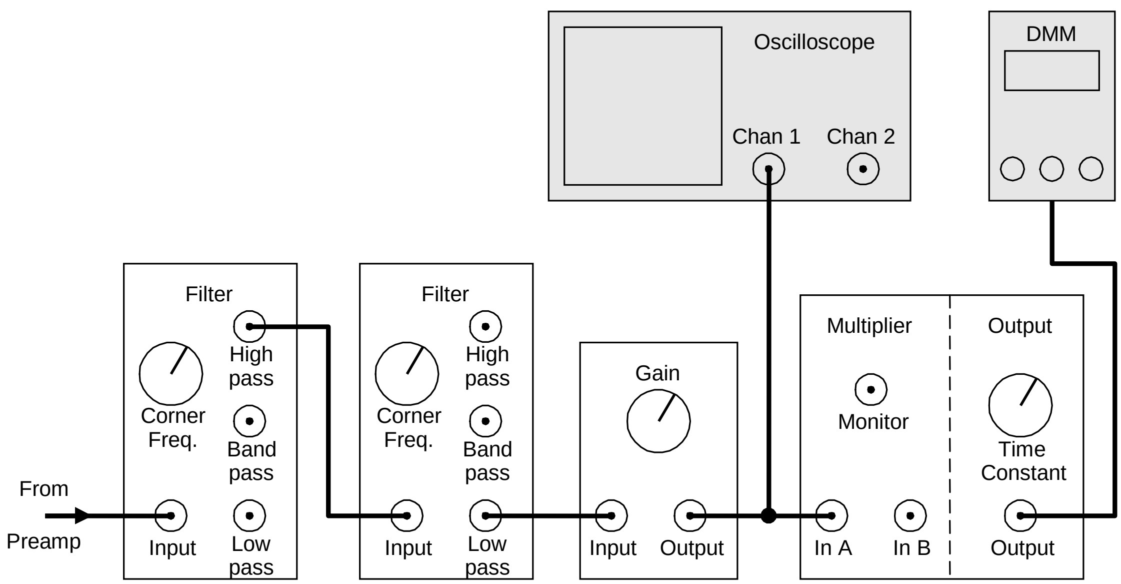

Fig. 1.2a: Cabling diagram for using the multiplier as squarer. High-pass filter 0.1 kHz, a.c. coupling; Low-pass filter 100kHz, a.c .coupling; Gain 400, a.c. coupling ; multiplier AxA, a.c. coupling

In fact, to persuade yourself that the squarer is working, use the XY-display capability on your 'scope. Convey the squarer's input $V_{in}(t)$, both to the squarer and to the X-channel of your 'scope, and convey $V_{out}(t)$ to the Y-channel, and have a look at a real-time XY-display. You should see a parabola emerge. See to it that you understand the origin of your XY-coordinate system, and then try changing some things: What are the right sensitivities to choose on the two axes? What would happen to your parabola if you raised the gain in the main-amplifier module of the HLE? Why does your data lie on a parabola, after all?

Now with*out* the need for a further cable, the output of the squarer is already being sent internally to the Meter module of your HLE. What this module does is to take the time-average of $V_{out}(t)$, averaged over a time interval you can select (by switch) to 1.0 second. This time average will *not* be zero, since $V_{out}(t)$, though fluctuating, is always and only on the *positive* side of zero. (Recall that the multiplier's squaring function ensures that $V_{out}(t)$ is proportional to the *square* of $V_{in}(t)$.) The meter will display that time-average, either on its 0-10 V or its 0-2 V scale. We suggest the use of the 0-2 V scale, and also suggest you go back and change the main-amp gain until the meter reaches a value near mid-scale, about 1 Volt on the 0-2 V scale.

What can you infer from this? Start with $V_J(t)$, the actual instantaneous Johnson-noise voltage generated by the source resistor. At the output of the pre-amp, you have a signal:

$$(6.00)(100.) \cdot V_J(t).$$

After the filter stages, you have the 0.1-100 kHz bandwidth-selected, or filtered, part of this signal. After the main amp, you have a signal

$$G_2 \cdot (600) \cdot V_J(t),$$

where G_2 is the main-amp gain, perhaps 300. Then after the squarer, you have a signal

$$[(300) \cdot (600) \cdot V_J(t)]^2 \,/\, (10\text{ V}) .$$

Finally, using the <...> brackets to indicate a time average, what you have displayed on your meter is the signal

$$V_{meter} = \langle V_J^2(t)\rangle \cdot (600 \cdot 300)^2/(10\text{ V}) .$$

From this result and the meter reading, you can work all the way backwards to find $\langle V_J^2(t)\rangle$, the mean-square voltage present (within your chosen bandwidth) across the source resistor.

Now use a cable to carry this time-averaged positive voltage to a digital multimeter. You should see a number consistent with your analog-meter indication, and you should see it fluctuate. (The expected size, and speed, of the fluctuations are treated in Appendix A.12.) Note that with the use of a 1-second time constant, you'll have to wait rather *longer* than one second for results to stabilize to any new value, especially if you're waiting for the 3rd or 4th digit of a multimeter display to settle down. Once the reading *has* settled, you'll notice the residual fluctuations, but go ahead and write down multiple readings from the multimeter, taking a new reading every second or so. See if you can persuade yourself that the readings display fluctuations about a mean value, and compute that mean value. It is connected, by a known chain of amplification and filtering, to the mean-square Johnson-noise voltage at the source.

1.3 Observing and Correcting for Amplifier Noise

You've now seen how all-analog electronics can take you all the way from a Johnson-noise source voltage $V_J(t)$ to a time-averaged d.c. voltage which is a traceable measure of $<V_J^2(t)>$. This section teaches you how to

a) make that measurement optimally, and
b) correct that measurement for amplifier noise.

a) The noise measurements you perform all depend on the linear operation of the amplifiers, and they (like all analog electronics) have only a ***finite range of output voltages over which they remain linear***. For the high level electronic amplifiers, that range is (-10 V, +10 V). If you were to put a simple sinusoid through the amplifiers, you could use the full ±10-V excursions . But since you are amplifying *noise*, you have to ensure that even the rare large fluctuations of the noise stay within the ±10-V 'span' of the amplifier. In practice, a maximum ***average noise signal*** of 3 Volts (rms) is a safe choice. This should avoid serious distortion of the signal, called 'clipping', like that shown in Figure 1.3a. For an average noise signal of 3 Volts rms, an excursion beyond ±10-V is so rare as not to spoil the accuracy of your measurement.

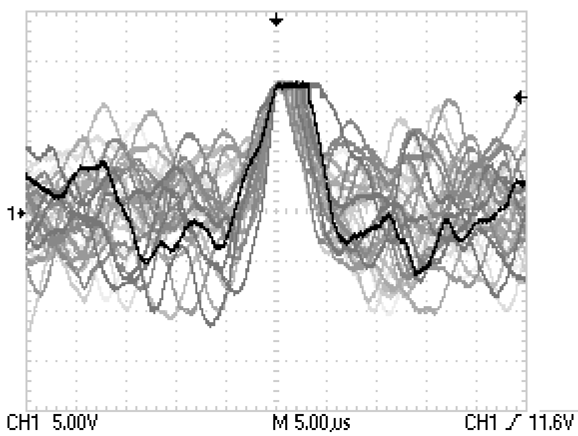

Figure 1.3a: Clipped signal from HLE – notice the clipping level is near +12 Volts.

Now if the rms measure of the signal at the A-input of the squarer, $V_A(t)$, is 3 V, then (by definition) its mean-square value is

$$<V_A^2(t)> = (3\text{ V})^2 = 9\text{ V}^2 ,$$

and under these circumstances, the squarer's MONITOR output will give

$$V_{sq}(t) = [V_A(t)]^2/(10\text{ V})$$

so that the time-average at the OUTPUT will be

$$<V_{sq}(t)> = <V_A^2(t)>/(10\text{ V}) = (9\text{ V}^2)/ (10\text{ V}) = 0.9\text{ V} .$$

You could use a smaller rms size for the input $V_A(t)$, but you'd be getting an even smaller output from the squarer, and your readings might be affected by zero-offsets in the squarer's output. (See Section 5.3 for details.)

So from here onwards, whenever you measure a noise voltage, you should check the main-amp output to see that it fits easily into the ±10-V range. If it exceeds these limits, reduce the gain. And you should look at the squarer's output on the panel meter, to see a time-averaged output near, or a bit below, 1 Volt. Again, if it's much larger, you want to reduce the gain, or if much smaller, raise the gain. Whenever you do take a reading of the time-average of the squarer's output, be sure to record also the net gain you've used to attain that reading, since this is your ticket to tracing the meter reading back to the desired mean-square noise $<V_J^2(t)>$.

b) Now back to Johnson noise. The problem you're now going to address is tracing noise back to a source, because here you have to consider the possibility that some of the noise you're seeing is *not* due to the Johnson noise of the of source resistor, but instead due to the amplifier chain which follows it. Since this 'amplifier noise' is just as featureless and random as the resistor's Johnson noise, there's apparently no way to separate the two waveforms *once they're added*. But there *is* a way to separate their effects, if we can assume that the amplifier noise does not depend on the source resistor's value. Here's the demonstration: let $V_J(t)$ be the instantaneous noise voltage from the source resistor, and let $V_N(t)$ be the instantaneous noise voltage apparently present at the input of the amplifier. That is to say, $V_N(t)$ is a model for a noise emf which, applied to the input of an ideal *noiseless* amplifier, would match the noise actually observed at the output of the real amplifier, driven only by its internal noise. If the gain of the amplifier is G, its output will be

$$V_{out}(t) = G\,[V_J(t) + V_N(t)]\ ,$$

and the mean-square of this output will be

$$<V_{out}^2(t)> = G^2 < [V_J(t) + V_N(t)]^2$$

$$= G^2\{<V_J^2(t)> + 2 <V_J(t) \cdot V_N(t)> + <V_N^2(t)>\}\ .$$

There's a 'cross term' in this expression, the time average of the product $V_J(t) \cdot V_N(t)$, but this time average is <u>zero</u>. The reason is that $V_J(t)$ and $V_N(t)$ can be safely assumed to be **<u>un</u>correlated**, arising as they do from distinct physical mechanisms in two different objects. So when $V_J(t)$ happens to be positive, the amplifier noise $V_N(t)$ is just as likely to be negative as it is positive; thus the product of the two factors is also as likely to be negative as positive. That's why the absence of correlation enforces a zero for the time-average of the product. But that fact leaves

$$<V_{out}^2(t)> = G^2\{<V_J^2(t)> + 0 + <V_N^2(t)>\}\ ,$$

which says that ***mean-square voltages from uncorrelated sources are simply additive***. In particular, it gives us a way to measure the amplifier noise -- we just change temporarily to a configuration in which the Johnson-noise term in this sum is negligible. Theory says that a choice of $R = 0$ for source resistance would give $<V_J^2(t)> = 0$, but in practice, it suffices to use the $R = 1\text{-}\Omega$ or $10\text{-}\Omega$ settings for giving a $<V_J^2(t)>$ which is small enough that the result is a good measure of the amplifier noise, $<V_N^2(t)>$.

And once that latter value is measured***, it can be assumed to be present, and unchanged, in any use of (the same configuration of) the amplifier***.[1] So for any source resistor R_{in} > 10 Ω, the amplifier noise contribution previously established can be subtracted off, leaving $<V_J^2(t)>$ isolated by itself.

Here's a concrete illustration: we have the values R_{in} = 1 Ω, 10 Ω, etc. We pick the 0.1 - 100 kHz bandwidth as before, and we pick gains to give good results at the squarer. In a particular example, the time-averaged outputs of the squarer we find are the $<V_{sq}>$ values below:

R_{in} chosen	gain G_2 (HLE)	$<V_{sq}>$ read	$<V_J^2 + V_N^2>$ inferred	$<V_J^2>$ derived
1 Ω	1500	0.6353 V	7.843 x 10^{-12} V^2	≈0.002 x 10^{-12} V^2
10 Ω	1500	0.6372	7.867	0.026
100 Ω	1500	0.6516	8.044	0.203
1 kΩ	1500	0.7911	9.767	1.926
10 kΩ	1000	0.9801	27.225	19.384

Now we expect, for the time-averaged output of the squarer,

$$<V_{sq}(t)> \quad = <V_{in}^2(t)> / (10\text{ V})$$

$$= \{(G_1\, G_2)^2 / (10\text{ V})\} <V_J^2 + V_N^2> ,$$

so we can use the G_1 = 600 and G_2-as-listed values to compute the column with $< V_J^2 + V_N^2>$ values. We can eyeball-extrapolate to the $R_{in} \rightarrow 0$ limit, and deduce a contribution of 7.841 x 10^{-12} V^2 for $<V_N^2>$ alone, the amplifier noise contribution (for this particular amplifier chip, at this particular bandwidth -- your number will vary!). Subtracting this contribution from all the entries gives the rightmost column for $< V_J^2 >$, our estimate of the mean-square Johnson noise of the source resistor, corrected for the effects of amplifier noise. Notice that the amplifier-noise corrections are large, even dominant, for small values of source resistance! You'll find (for the present choice of pre-amp input stage) that Johnson noise surpasses amplifier noise only when the source resistance has risen to about 3 kΩ.

[1] Under the assumption of negligible op-amp current noise, and no noise from external interference, both of which may depend on $R_{in.}$

1.4 Johnson noise dependence on resistance

The previous sections have taught you how to configure the pre-amp/filter/main-amp combination, and how to select a gain for optimal use of the squarer. The results can also be corrected for amplifier noise, and traced back to an inferred mean-square measure of Johnson noise, $< V_J^2(t) >$, for any source resistor from R = 10 Ω upwards.

You should now investigate systematically the dependence of $< V_J^2(t) >$ upon source resistance R. To do so, you can use the R = 10 Ω through 10 MΩ choices built into the pre-amp module. (These internal source resistors have tolerances of 0.1% to 1 MΩ, and 1% thereafter.)

In addition to the measurement using the fixed internal resistors, you can also measure arbitrary resistors. The selector switch gives you access to three more test positions, A_{ext}, B_{ext}, and C_{ext}. You can open the case to populate these external positions with resistors as you please. Measurement of these resistors is optional, so you don't need to do so in order to finish the practicum. If you want to do it, please contact your advisor before opening the panel.

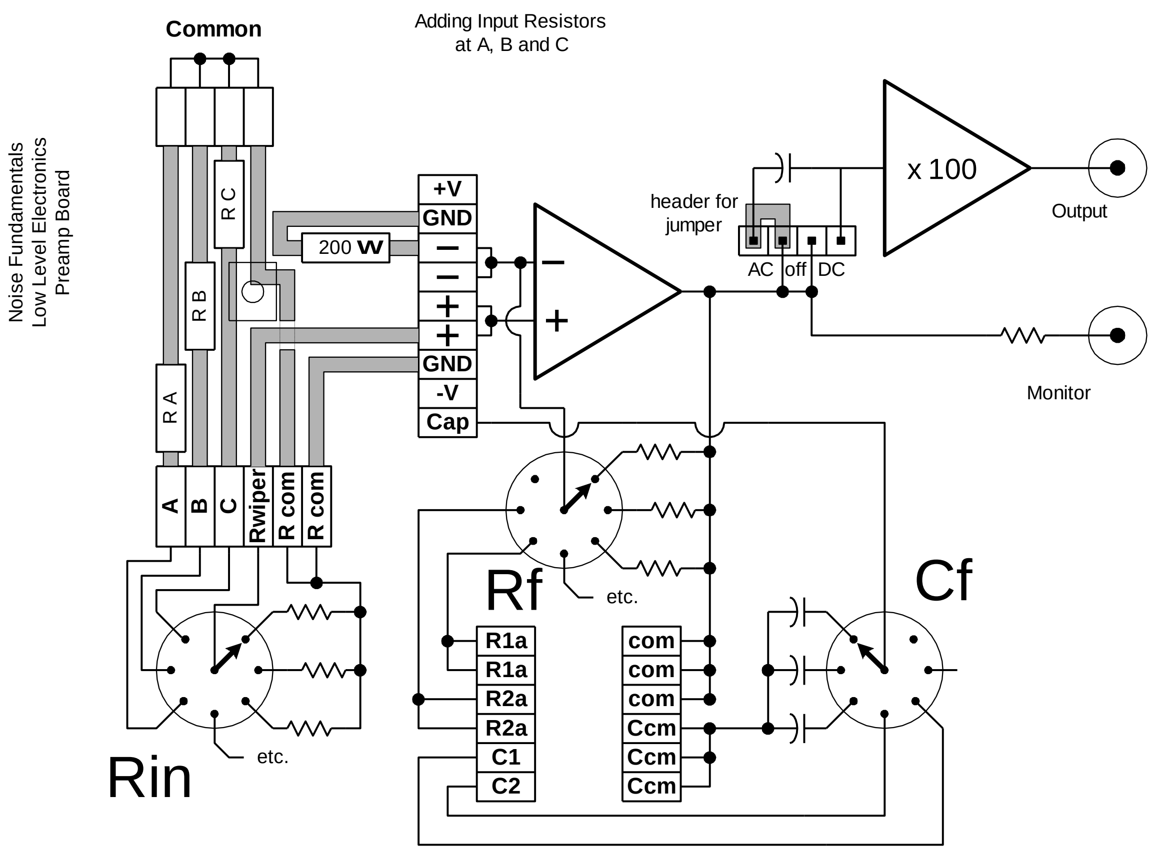

Fig. 1.4: Wiring diagram for adding components at the A, B, C, positions of the pre-amp's input. Note all input resistors have a common ground.

After your advisor gives the OK, you can change the resistors by doing as described in the following: First you can 'flip' the pre-amp panel, as illustrated in Figure 7.2a, to expose the back (component) side of the pre-amp's circuit board. You can also find the pre-amp power switch (near the internal power-on red LED inside the low-level electronics), and **turn OFF the pre-amp power** before making any changes to the board. Now use the diagram below to find the screw-connect terminal strips, and find also the location of the two endpoints for the components you're putting into the A, B, and C positions.

You can choose resistors of any value in the 20 Ω to 5 MΩ range; you can even choose different kinds of resistors. (Most resistors sold nowadays are of metal-film construction, but ask around for some carbon-composition or wirewound resistors -- and look up what kinds of resistors Johnson himself used.) You can clearly use resistors of any power capability you like -- their internal Johnson-noise emf is *not* going to overheat them! If you wish, you can have a comrade *hide* from you the resistance values, so you'll be measuring some actual unknowns. Don't forget to turn the pre-amp power back ON before you re-flip the front panel and close up the box.

Now you can take noise data for your own resistors, as well as for the built-in source resistors. If you didn’t build in external resistors restrict your measurements on the internal resistors. Once you have values for $< V_J^2\,(t) >$, each corrected for amplifier noise, you can plot those values as a function of R. Since both axes will vary over many orders of magnitude, a log-log plot is appropriate. The vertical axis has units of Volts-squared, the horizontal axis has units of Ohms. Nyquist's theory predicts a first-power power-law dependence on resistance R, namely

$$<V_J^2(t)> = (4\, k_B\, T\, \Delta f) \cdot R^1\ ,$$

and you might see this confirmed. There will be deviations from this behavior at the high-R end of the plot, for reasons to be discussed in sections 1.5 and 2.2, and Appendix A.8.

At the *low*-resistance end of the plot, you'll see the amplifier-noise-corrected values enable you to follow Johnson noise to a regime *well* below the apparent limit set by amplifier noise. You'll be able to establish values of $<V_J^2(t)>$ which are less than 1% of the amplifier noise $<V_N^2(t)>$ that overlays them. Of course, the corrected value of Johnson noise will be the difference between two nearly equal quantities, so the results will be subject to larger uncertainties than other data points.

1.5 Johnson noise dependence on bandwidth

Thus far you've learned how to observe and quantify Johnson noise, and you've seen how to isolate its mean-square value from amplifier noise. You've also seen its dependence on source resistance R. But Nyquist's formula claims that $<V_J^2(t)>$ also depends on the bandwidth Δf; ie. on the range of frequencies to which your system is sensitive. [2]

So for now you should stay at room temperature, and stay at a fixed R-value; we suggest a starting value of R_{in} = 10 kΩ. The goal is to see how the choice of bandwidth matters. The method is to imagine a 'white noise spectrum', ie. noise power uniformly spread in frequency at its origin, but subsequently modified by the high-pass and low-pass filter sections as depicted below.[3]

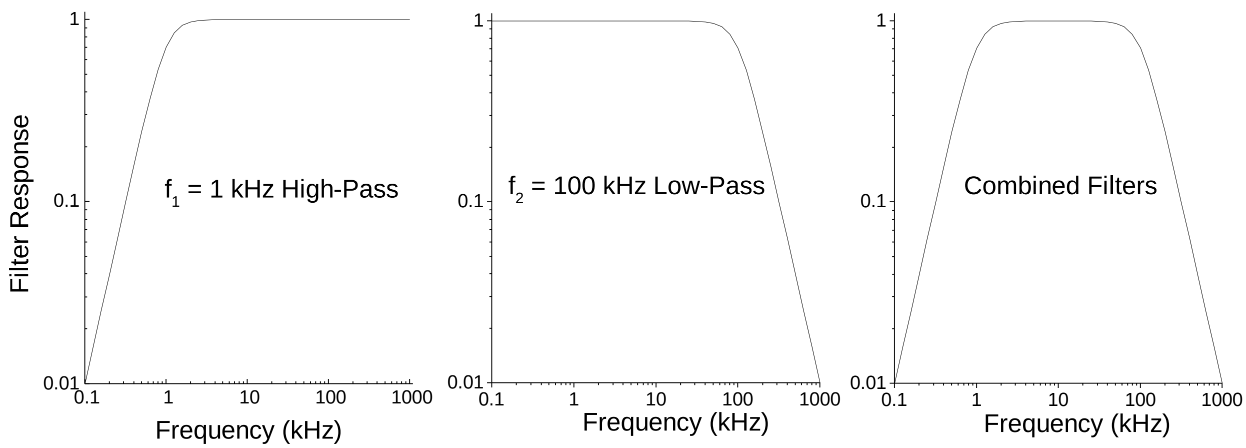

Fig. 1.5a: Representation (left) of the transmission of a high-pass filter, of corner frequency f_1; (center) of a low-pass filter, of corner frequency f_2; (right) the combined effect of both filters. The graph's scales, horizontal and vertical, are all logarithmic.

You have a range of choices for the 'lower corner' frequency f_1 or high-pass filter setting, and a separate range of choices for the 'upper corner' frequency f_2 or low-pass filter setting. You might first think that the bandwidth Δf should be given by $|f_2 - f_1|$, which is a decent approximation, but subject to significant corrections. These so-called 'corrections' are discussed in great detail in Sections 2 and 5. But for now we present you with the generic corrections which are the result of a model calculation. The model of Section 2.2 predicts the effective bandwidth Δf for each combination of f_1 and f_2, and gives the results shown in Table 1.5.

---

[2] The further prediction that Johnson noise depends on the resistor's *temperature* is tested in Chapter 4.

[3] Section 2.0 teaches you how to get data of this form.

Table 1.5 Effective noise bandwidths, Δf, given in Hertz, computed for model filter responses

	f_2 = 0.33 kHz	1 kHz	3.3 kHz	10 kHz	33 kHz	100 kHz
f_1 = 10 Hz	355	1,100	3,654	11,096	36,643	111,061
30 Hz	333	1,077	3,632	11,074	36,620	111,039
100 Hz	258	1,000	3,554	10,996	36,543	110,961
300 Hz	105	784	3,332	10,774	36,321	110,739
1000 Hz	9	278	2,576	9,997	35,543	109,961
3000 Hz	0.4	28	1,051	7,839	33,324	107,740

These computed values are all subject to uncertainties of order 4%; (see Section 5.2 for details on how any of them can be more carefully calibrated). They are all computed (by the methods of Section 2.2) for ideal filter responses, ignoring systematic effects. Inclusion of those effects may raise values in the rightmost column by (3±1)%, and may raise values in the next-to-rightmost column by (1±1)%. There are further corrections to effective noise bandwidths for large f_2-values, in the case of large source resistance, due to capacitive effects -- see Appendix A.8.

Your goal is to measure the mean-square Johnson noise of the resistor, $<V_J^2(t)>$, for as many (f_1, f_2) combinations as you wish. Recall that for each choice of filter settings, you'll want to adjust the gain so as to use the squarer optimally. Recall that each mean-square value you measure needs to be corrected for amplifier noise (measured at *that* bandwidth setting: the amplifier-noise contribution to the mean-square depends, as does the Johnson-noise contribution, on the bandwidth you use.)

You can plot your data for $< V_J^2(t) >$ in various ways:

- as a function of the f_1-value used to obtain it;
- as a function of the f_2-value used to obtain it;
- as a function of the difference $|f_2 - f_1|$ of the f_1- and f_2-values used to obtain it; or
- as a function of the equivalent noise bandwidth, from the table above.

Which plot is the most nearly linear? Try again using log-log scales, to be able to see all you data points, spread as they are over many orders of magnitude. Pick the most informative plots to include into the protocol.

If your plot is consistent with $< V_J^2(t) > \propto \Delta f$, then the coefficient of this proportionality tells you a 'noise power spectral density', as you'll see in the next Section. Its units are V^2/Hz, and it's usually denoted by S.

1.6 Johnson noise density, and Boltzmann's constant

Previous sections have shown you how to measure noise, and have tested its dependence on source resistance R and on measurement bandwidth Δf. This section introduces you to noise *density*, and then relates your measured values, via Nyquist's formula, to Boltzmann's constant.

If you have shown that measured mean-square noise $\langle V_J^2(t)\rangle$ has a linear dependence on the bandwidth Δf used, you are entitled to infer the existence of a 'noise density' that's uniform in frequency. Here's an analogy to <u>mass</u> density that should make this clear -- we'll use a one-dimensional example. Suppose you have a string, of unknown composition, laid out on an x-axis, and that you can make clean cuts at arbitrary locations x_1 and x_2, and then weigh the piece of string you've extracted. If (and only if) you find that the observed mass M is always proportional to $|x_2 - x_1|$, you may conclude the string is of uniform density. You can also see that the quotient

$$(\text{mass } M) \,/\, |x_2 - x_1|$$

gives the value for this density, given in units of mass per unit length.

Similarly, if you've shown that mean-square noise $\langle V_J^2(t)\rangle$ is always proportional to the bandwidth Δf you used to obtain it, then you can define the 'noise power density'

$$\langle V_J^2(t)\rangle \;/\; \Delta f \,,$$

in this case with units of Volts-squared per Hertz, or V^2/Hz. [Strictly speaking, this is not a power density -- but if a voltage $V(t)$ is applied across a resistance R then the quotient $V^2(t)/R$ <u>is</u> a power. So the quotient above is just a factor-of-R away from being an actual power density, with units Watts per Hertz.]

Your data for a single source resistance R = 10 kΩ has given you a noise power density; you can go back to your data of Section 1.4 and convert that data to noise power density as well, to check the dependence-on-R of this density. The motivation for all of this is that Nyquist's formula can be written as

$$\text{noise density } S = \langle V_J^2(t)\rangle \,/\, \Delta f = 4\, k_B\, T\, R \,.$$

So you should plot all of your data thus far, and perhaps more data that you now take for various R- and Δf -values, to see if you can further establish the linear-in-R claim of the prediction above. (In practice, you'll see deviations in the regime where R and/or Δf is large, for reasons discussed in Appendix A.8.)

If you establish a regime of linear dependence on R, your plot, or fit, will give you a value for a slope, $(4\, k_B\, T)$. What *units* should it have? (Answer: rise over run, so V^2/Hz per Ohm -- and what unit is *that*?) What *value* does it have? Hardest: what *uncertainty* can you assign to your value? (Do so <u>before</u> you look up any 'book values', because the uncertainty intrinsic to your experiment is conceptually a matter quite separate from any discrepancy between your value and anyone else's.)

Finally, if you know you room's temperature T (and express it in absolute, ie. Kelvin, units), you can now conclude by finding a value (and uncertainty!) for Boltzmann's constant k_B.

What's the nature of k_B? At one level, it connects historical choices of temperature units to a 'common language' of energy units, via $E = k_B\ T$. At another level, k_B is a 'microscopic' version of the macroscopic gas constant R (here, not a resistance), as you can see by writing the ideal-gas law in two ways,

$$p\ V = n\ R\ T \qquad \text{and} \qquad p\ V = N\ k_B\ T\ .$$

The first form has n = (number of moles of gas), and that gives to R the units of Joules per (mole·Kelvin). The second form has N = (number of molecules of gas), and it gives to k_B the units of Joules per (molecule·Kelvin), or just J/K. This double form of the law also makes it clear how R and k_B have to be related: since ($n\ R$) and ($N\ k_B$) both give pV/T, we have

$$n\ R = N\ k_B\ , \qquad \text{or} \qquad R = (N\ /\ n)\ k_B\ .$$

But (N/n) is Avogadro's number N_A, the number of molecules per mole. Hence you expect the numbers to obey the relation

$$R \approx 8.31\ \text{J/mole·K} = N_A \cdot k_B \approx (6.02 \text{ x } 10^{23}\ \text{/mole})\ (1.38 \text{ x } 10^{-23}\ \text{J/K})\ .$$

Check that claim. Does this mean that electrons inside a resistor are acting like molecules in a gas, bouncing around between the resistor's two ends? And is the Johnson-noise emf akin to the pressure fluctuations which kinetic theory predicts for a gas? What's the connection to Brownian motion? See if you can find any guidance on these conceptual points.

.

2. Noise as a function of temperature

2.1 Equipment, methods, and issues

This section explains the use of the 'thermal probe' and its associated Dewar vessel. Together they make it possible to measure Johnson noise from a source resistor as a function of its temperature T. The system is designed for use with liquid nitrogen (LN_2) as a coolant, and an electrical heater allows exploration above that base temperature. The probe is suited for use in the 77 K - 400 K range. The lower end is set by the normal boiling point of LN_2; the upper end (which is +127 °C) is set by temperature limits of wires and components in the probe head, and is enforced by the limited power available to the heater.
You can measure the room temperature values on your own. After that, please contact your advisor, who will provide you with the liquid nitrogen.

The motivation for this temperature coverage is of course the theoretically-predicted ($4\ k_B T\ R\ \Delta f$) **temperature** dependence of the mean-square Johnson noise voltage. Using the accessible temperature range, you'll be able to vary this quantity by a factor of 4 or 5.

SAFETY WARNING: The Dewar supplied is made of un-silvered glass to help you see the liquid level inside. Because it's made of glass, it will shatter if you drop it. The disaster will be even more dangerous if the Dewar is full of LN_2 when dropped.
So: *do **NOT*** drop the Dewar, and use and store it only in the base built to hold it securely.

SAFETY WARNING: Liquid nitrogen is *very* cold, boiling at about -195 °C. It is dangerous to have it contact your skin, and even more dangerous to undergo skin contact with clothing soaked with LN_2. The hazard is not chemical, but physical. You can suffer frostbite, and permanent nerve and/or tissue damage, from the localized freezing that will occur.

There is also a special electronics issue involved with the use of the probe. The source resistors are now not built into the pre-amp module, but instead a few feet away. They are connected to the first stage of amplification by wires inside the low-level electronics box, and then by a coaxial cable over to, and down into, the probe. We have succeeded in preserving the required electrical grounding and shielding of those remote resistors against external electrical interference; but the inevitable cost is much larger capacitance between the 'live wires' and the shields. This capacitance (about 100 pF) has consequences on the bandwidth of the noise signals. *The Johnson noise is still spectrally 'white'* ***at its origin****, but its spectrum is already modified by capacitive effects when it reaches the first amplification stage*. Some solutions to the problem will be presented.

2.2 Two-temperature Johnson-noise measurement

This section teaches you how to prepare your system for measuring Johnson noise in 'remote resistors'. It pre-supposes that you've worked through Chapter 7 and Sections 1.1-1.5 on how to measure Johnson noise in 'local resistors'. The goal is to measure Johnson noise at two distinct temperatures: ambient and liquid-nitrogen.

The first thing you'll need to do is to confirm the installation of resistors into the probe. The unit is shipped with resistors R_A = 10 Ω, R_B = 10 kΩ, and R_C = 100 kΩ already installed, as you can see in the Figure 4.2a:

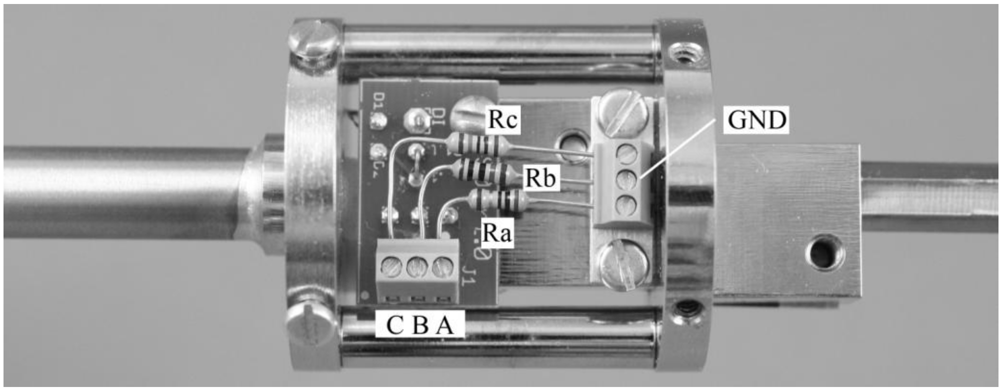

Fig. 4.2a: The interior of the temperature probe, showing the A, B, and C positions of source resistors.

Notice that to get this view, you have to loosen four screws, and remove four more, to slide away the shielding sleeve on the probe. The resistors' 'hot ends' are on the circuit board, and the ends near the copper flange are grounded. Once you've confirmed the resistors are present, you need to close up the shielding sleeve again.

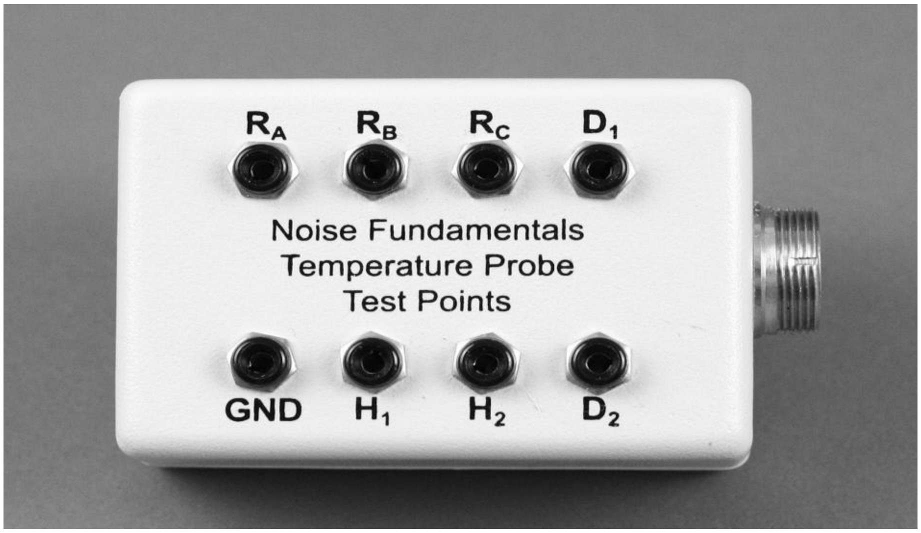

Figure 4.2b The 'breakout box'

To facilitate checking that all the components are properly connected inside the variable temperature probe, we have included with the unit a 'breakout box' shown in Figure 4.2b. This box connects to the variable-temperature probe and has 8 test points, one for each wire lead from the components in the probe to the connector. Points R_A, R_B, R_C, and GND can be used to test the resistors with an ohmmeter, since all three resistors have a common ground. (Note that the $10\,\Omega$ resistor will likely read $12\,\Omega$ because of the resistance of the leads.)

The heater leads are present at H_1 and H_2 which measure about $75\,\Omega$. The diode thermometer should be checked with a multimeter (on the diode-testing scale). It should read about half a volt, if the positive lead is connecter to either D_1 or D_2 . Note that there is only one diode thermometer connected at the factory and *both* wires D_1 and D_2 are connected to it.

Internal wiring brings three 'live wires' from the three source resistors to a junction box atop the probe, and then via a cable to a connector for the Temperature module of your low-level electronics. Before you connect the probe to that module, open up the low-level electronics (by the familiar flip operation) to see what connections you need to make between the Temperature module and the Pre-amp module. The Figure 4.2c shows the necessary connections. In particular, you need three wires connecting the R_A, R_B, and R_C resistors to the A_{ext}, B_{ext}, and C_{ext} positions in the pre-amp. The 'other wire' of each of the three resistors is already grounded to the body of the probe, as the resistors are both thermally and electrically connected to the copper block through the terminal post .

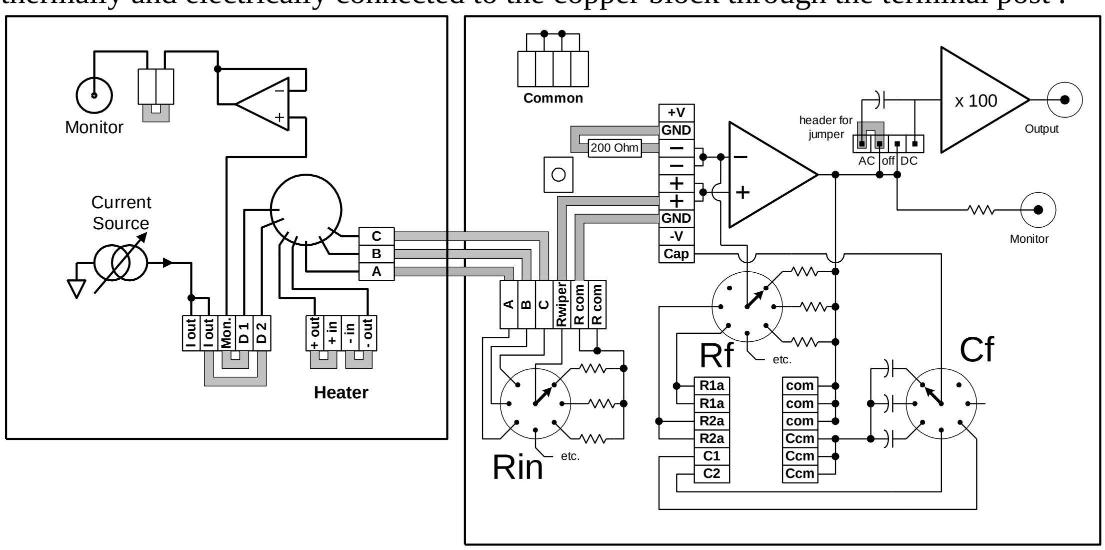

Fig. 4.2c: Wiring diagram of the interior of the low-level electronics, to bring the A, B, and C remote resistors to the A_{ext}, B_{ext}, and C_{ext} positions of the R_{in} selector, and to connect the temperature transducer and heater.

Also shown in the diagram are the connections you'll want to make for the temperature transducer, and the heater, in the probe. You'll need those devices in future sections.

Re-flip the low-level electronics into its enclosure, confirm its power is ON, close up the box, connect the probe cable to the thermal module, and install the probe as shown in the photo below. Note that the Dewar is absent, and the whole probe is at ambient temperature.

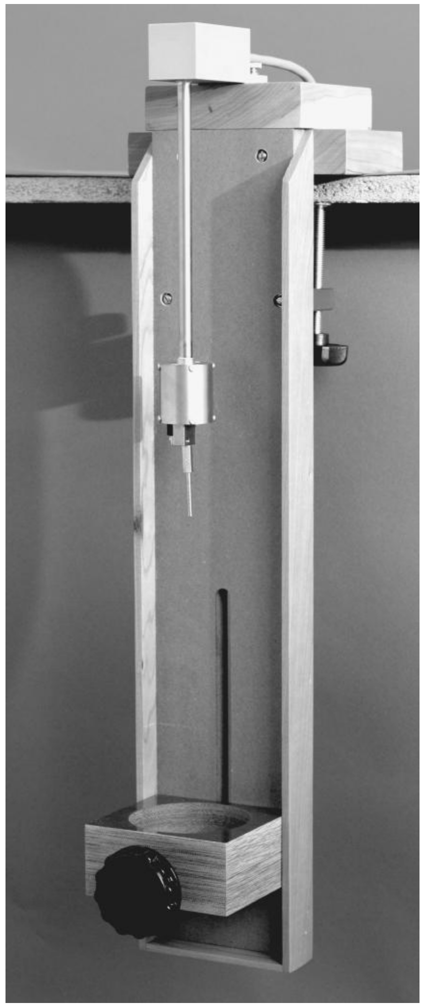

Fig. 4.2d: A mounting for the Dewar support, and temperature probe (with Dewar vessel removed).

Now you should be able to measure (room-temperature) Johnson noise from three remote resistors, just by using the A, B, C, positions of the source-selector (R_{in}) switch on the pre-amp. Other positions of this switch make available the noise from a set of 'local' resistors, including values of 10 Ω, 10 kΩ, and 100 kΩ.

For initial measurements, we suggest a bandwidth of about 10 kHz (set perhaps by using a 1-kHz high-pass, and a 10 kHz low-pass, filter). As usual, you'll need to recall that the standard gain is G_1 = 600 in the pre-amp (if that's in its default condition), and you'll need to use a suitable gain G_2 in the main-amp to get the squarer to operate in its optimal regime. As previously, here too you'll need to use the 10-Ω source resistor as a way to get the amplifier-noise contribution, which needs to be subtracted from the mean-square noise measurements.

It is important that you take data from both local and remote 10 kΩ and 100 kΩ source resistors, and also that you try some different bandwidths. Because of the effects of probe capacitance, it is to be expected that the values of $\langle V_J^2(t) \rangle$ you infer will be smaller for the remote, as compared to the local, resistors. The deficiency will be larger for the larger source resistance, and a broader bandwidth. To see why this is expected, compute an RC time-constant for choices of R of 10 kΩ and 100 kΩ, now assuming $C \approx 100$ pF for connections to the probe resistors. Then compute a corner frequency of the undesired one-pole low-pass filter that results, from $f_c = 1 / (2\pi\ \tau)$. See Appendix A.8 for how to handle the consequences.

Use these (all room-temperature) results to decide on a measurement strategy that you'll use when the probe is *not* at room temperature. When you've worked that out, it is finally time to cool your probe. The photo above suggests how a (warm, and empty) Dewar can be slid into place, mounted into its movable base, and used to surround the probe. You can lower that base and the Dewar together and pour about 1 liter of LN_2 into the Dewar. [Go back to Section 4.1 and re-read the SAFETY WARNINGS we've posted there -- liquid nitrogen is a tool, or a hazard, **but not a toy**.] Wait for the boiling to subside, slide the black foam insulating cover down onto the Dewar's mouth, and now use the clamp on the Dewar's base to raise the Dewar until the probe makes contact with the LN_2. Here, as in general, the probe's sample chamber should end up at about the mid-height in the Dewar, and (for purposes of *this* experiment, exceptionally) also to end up with its copper bottom plate immersed in the liquid. (In later sections, you'll want only the brass 'cold-finger' on the bottom of the probe to be immersed.) The purpose is to ensure that your resistors really are at the temperature of your boiling LN_2.

When all the extra boiling has settled down, you can repeat your noise measurements, using both the local and the remote resistors, and using the protocol you've established. You may need to change the gain G_2 to keep the squarer in its optimal regime.

There's one more necessary measurement task. Your 'remote resistors' are of 1% tolerance, but that does *not* guarantee that their 77-K resistance matches their nominal value to this accuracy. So you'll want to check their 'cold resistance', ie. their R-values when they're immersed in LN_2. To get access to their electrical properties, we've supplied a 'breakout box', to which you can attach the cable of the probe, to get connections with all the items down inside it (see Appendix A.1). For remote resistors immersed in freely-boiling LN_2, you can not only be pretty sure of their temperature, you can also be quite confident that the diagnostic currents used by an ordinary ohmmeter will not warm the resistors significantly, as you measure their values.

Here's a final note -- you might get to this point, and for the first time have a cold probe immersed in leftover LN_2. Here are some suggestions for what to do at the end of a day's experimentation.

When you're done with your work, it might be a good idea to lower the Dewar's base, remove the Dewar, dispose of the surplus LN_2 by your locally acceptable method, and lay the Dewar down on its side to warm up. (Why is this better than leaving it standing vertically? If you do lay it down, do NOT let it roll away to its doom.) This removal will leave a cold probe in the open air, and you do NOT want to touch it -- contact with cold metal can lead to immediate frostbite, as well as the dreaded 'pump-handle effect'. Instead, leave the probe to hang in ambient air, and warm up spontaneously. If you're in a hurry, or if condensation of water onto the chilly probe is a problem, leave it in open air, with the heater running, perhaps set to 5 or 6 turns on the dial. Then it will warm and eventually equilibrate to a safe-to-touch temperature, yet far enough above ambient temperature to ensure that it dries out properly.

Historical insight: Once you have data of the sort acquired above, here's one use you can make of your ambient-temperature and LN_2-temperature values for $\langle V_J^2(t) \rangle$. Put yourself back into the era of the Centigrade scale of temperature, on which ice melted at 0 °C and water boiled at 100 °C, by definition. On such a scale, your ambient temperature might be 22 °C, and your LN_2 temperature might be -195 °C (find an old reference which quotes you this value -- how do you suppose that it was established?). Now plot your two $\langle V_J^2(t) \rangle$ points as a function of the Centigrade temperatures at which they were measured. You have only two points, so of course you can fit a line to the two points. The pay-off is to find the $\underline{x}$-axis intercept of this line, as the extrapolated low-temperature point at which Johnson noise vanishes.

What you're doing is 'locating absolute zero' according to a noise-based measurement. It is a non-trivial technical, and intellectual, challenge to test whether Johnson noise extrapolates to zero at the same temperature at which the pressure of an ideal gas extrapolates to zero. Success in such tests suggests that the Kelvin scale is not just absolute, but also physics-wide. From a modern point of view, we depend on such a result to enable us to claim that the *T*-variable which appears in the Johnson-noise equation really is the absolute temperature, ie. the temperature measured relative to the absolute zero which is established by procedures such as these.

Appendices

Appendix A.1. Technical specifications

Low-level Electronics:

Pre-amplifier module:

First stage is user-configurable (see Appendix A.4.)

As shipped, first stage is

FET-input operational amplifier, non-inverting mode

gain (using R_f = 1. kΩ) is (1 + R_f /200. Ω) = 6.00

-3 dB bandwidth > 1.0 MHz

input impedance > 100 MΩ

Next stage are fixed-configuration

gain 100.

-3 dB bandwidth > 1.6 MHz

Temperature module

Current source

accuracy <1%, 10 nA to 1 mA settings

Transducer voltage buffer

gain 1.00 to <0.1% error, ±2 mV d.c. offset

Heater power supply

0 - 25 V (for floating loads), 330 mA current capability

Signal Attenuator

input impedance: variable, 100 Ω to 10 kΩ

output impedance: 10 Ω (for use driving Z_{in} = 1 kΩ stages)
or 10 Ω less 1% (for driving Z_{in} ≥ 10 kΩ stages)

-3 dB bandwidth > 10 MHz

Bipolar Power Supplies

output: (±) 10 mV to 11 Volts

noise: < 5 nV/√Hz, typically < 2 nV/√Hz

current capability: 250 mA maximum

High-level Electronics:

Filter sections
- state-variable 2-pole Butterworth design
- input impedance 10 kΩ

Main amplifier
- two stages, of gain x1 or x10 selectable
- one stage, gain variable from x10 to x100
- -3 dB bandwidth > 1.4 MHz
- slew rate ≈ 20 V/μs
- input impedance 1 kΩ

Multiplier
- scaling factor for output: $V_{out} = V_A V_B / (10.0\ V)$
- input impedances of A and B channels: 50 kΩ
- d.c. offset: under ±10 mV

Output stage
- hard-wired d.c. coupling to output of multiplier
- two successive (buffered) stages of 1-pole, low-pass filters

(back panel) Noise Calibrator
- output level ≈ 212 ± 2 mV, rms measure
- noise power is located >99% in $0 < f < 32$ kHz range
- spectral density uniform to ± 2 % in $0 < f < 32$ kHz range

The 'Break-out Box' for the Thermal Probe:

In normal operation, the cable from the Thermal Probe into the Temperature-Control module connects all the devices in the probe to circuits in the module. It does so in a shielded, all-grounded, low-noise environment. But there are times when you want ordinary access to connections in the probe -- for example, if you want to use an ohmmeter to diagnose the resistance (at ambient, or at LN_2, temperatures) of the resistors mounted in the probe. To do that, you can disconnect the Probe's cable from the Temperature-Control module, and connect it instead to the plastic breakout box. Now you won't have full shielding, but you *will* have test-probe access to wires:

GND	labels ground, ie. the shell and body of the probe, including the copper fin at its bottom
R_A, R_B, R_C	label the three source resistors' live ends (each has its other end grounded)
D_1, D_2	label the two wires from the temperature-monitoring transdiode (see Section 4.3); this transistor has its collector and base leads grounded, so the 'live wires' D_1 and D_2 connect to the emitter
H_1, H_2	label the two ends of the 75-Ω heater on the lower fin of the probe; this resistor has *neither* end grounded

Appendix A.2. The matter of a.c. or d.c. coupling

Real amplifiers are subject to 'd.c. offsets', such that a potential difference of zero at the input can still lead to a non-zero steady d.c. value at the output. Because the overall gain of the Noise Fundamentals system can be as high as (600) x (10^4) = 6 x 10^6, even an effective 1 μV offset at the pre-amplifier's input stage would lead to a full 6-Volt offset at the main amplifier's output. The amplified noise voltages would be lying atop that d.c. offset, and this would create unacceptable errors. So at many stages of the electronic signal chain, there is the option to use a.c. coupling between the stages.

Every a.c.-coupled connection (including that selection at the input of test instruments) is actually a high-pass filter, with a corner frequency typically located at 10 Hz or so. Thus the d.c. component of any signal is entirely blocked, and high-frequency a.c. signals are entirely passed, by the filter. But a.c. signals of frequencies below 10 Hz can be considerably attenuated, as well as phase-shifted, by the filter in question. This attenuation *matters* if the study of low-frequency noise is of interest to you.

What follows is a description of the a.c. vs. d.c. coupling options, stage by stage, in the Noise Fundamentals signal path.

The pre-amplifier's first stage is always d.c. coupled, as that's a necessity in shot-noise measurements. A MONITOR output allows a view of the d.c. output level of the first stage; any a.c. signal or noise is lying super-imposed on that d.c. level.

The gain-100 stage in the pre-amp can be a.c.- or d.c.-coupled to the first stage's output. As shipped, the coupling is a.c., with a high-pass corner at frequency 16 Hz. The change to d.c. coupling can be made via a moveable jumper on the printed-circuit board inside the pre-amp.

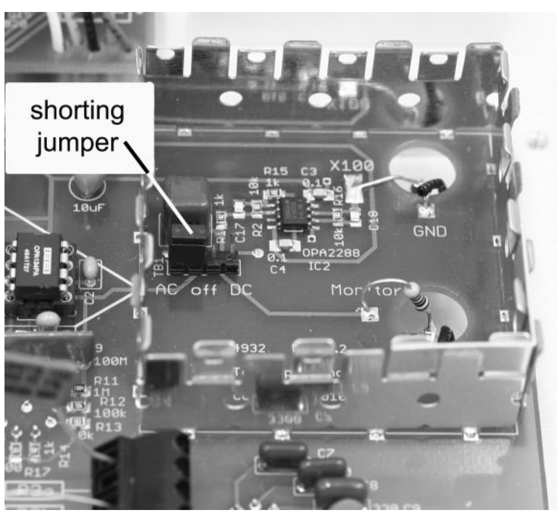

Fig. A.2a: How to select a.c. vs. d.c. coupling between the input stage, and the gain-of-100 stage, of the pre-amp.

In the high-level electronics, the two filter sections can be a.c.- or d.c.-coupled by front-panel switches. In the a.c.-coupled mode, there's a high-pass corner at frequency 1.6 Hz. In the main-amplifier section, the input can be a.c.- or d.c.-coupled by front-panel switch. In the a.c.-coupled mode, there's a high-pass corner at frequency 16 Hz.

The multiplier's inputs can both be grounded, or configured with a.c. or d.c. coupling. In the a.c.-coupled mode, there's a high-pass corner at frequency 1.0 Hz.

The output stage is internally connected, by d.c. coupling, to the multiplier's output, and its averaging action is optimized for accuracy all the way down to d.c.

The Noise Calibrator output is d.c. coupled, to preserve the flatness of its noise spectrum down near zero frequency. As a result, there may be a milliVolt-level d.c. offset in its average value.

Finally, a word about the consequences of a d.c. offset on an a.c. noise voltage going into the squarer. Recall that in typical operation, gains are chosen so that the signal reaching the squarer has an rms measure of about 3 Volts. Suppose that the actual signal entering the squarer is

$$V_A(t) = D + \Sigma_i A_i \cos(2\pi f_i t + \phi_i) ,$$

where here D represents the d.c. offset, and the sum is a Fourier representation of all the component frequencies in the signal (or noise). Since the squarer gives output

$$V_{sq}(t) = [V_A(t)]^2/(10\text{ V}),$$

the instantaneous output of the squarer contains lots of terms:

$$\begin{aligned} V_{sq}(t) = (10\text{ V})^{-1} \{ & D^2 + \Sigma_i A_i^2 \cos^2(2\pi f_i t + \phi_i) \\ & + \Sigma_i D A_i \cos(2\pi f_i t + \phi_i) \\ & + \Sigma_{i,j} A_i A_j \cos(2\pi f_i t + \phi_i) \cos(2\pi f_j t + \phi_j) \}. \end{aligned}$$

Upon taking the time average, the terms in the last two lines average to zero, while the cosine-squared terms average to 1/2. So what you observe as the time-averaged output is

$$\langle V_{sq}(t)\rangle = (10\text{ V})^{-1} \{ D^2 + \Sigma_i A_i^2 (1/2) + 0 + 0 \}.$$

The result is that the expected and intended output,

$$\langle V_{sq}(t)\rangle = \langle [\text{a.c. part of } V_A(t)]^2 \rangle / (10\text{ V}) ,$$

is polluted by an error of

$$\delta\langle V_{sq}(t)\rangle = D^2 / (10\text{ V}) .$$

So if the output of the main amplifier has an offset of even 100 mV, lying underneath the typically 3-Volt (rms) noise signal, and if the squarer is used in its d.c.-coupled mode at the A-input, this error will be $(0.1\text{ V})^2/(10\text{ V}) = 0.001$ V, relative to an output due to the intended noise signal of $(3\text{ V})^2/(10\text{ V}) = 0.9$ V.

This offset of 1 part in 900, or 0.1%, will not be caught or corrected by switching the input of the squarer to the ground (GND) position, since in this position the squarer does not get to see the d.c. component that might be present in the main-amp's output.

The moral of this story: unless you have reason or need to study noise below about 20 Hz, use a.c. coupling throughout. If you do use d.c. coupling at various places in the system, monitor the signal (being sure to use an oscilloscope set for d.c. coupling at *its* input!) at every point in the signal chain, to ensure

- that nowhere is the d.c. level sufficient to saturate the next stage, and
- that the d.c. average level underlying the input to the squarer is under 100 mV or thereabouts.

Appendix A.3. Operational-amplifier circuits and noise

This Section describes how the amplifier noise in operational amplifiers can be modeled, and goes on to discuss the implications for experimentation.

We start with the open-loop model for a (bare) op-amp, with two inputs (called inverting and non-inverting, but labeled by - and + respectively):

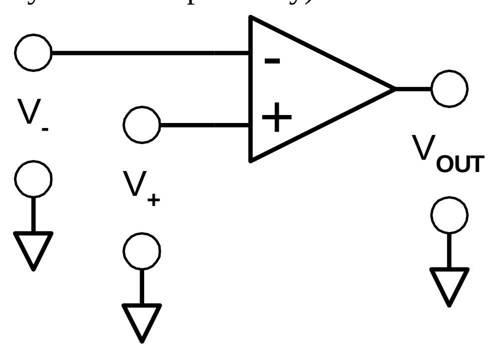

Fig. A.3a: An operational amplifier without feedback, showing labeling of inputs.

The noise-free model behavior is $V_{out} = A\,(V_+ - V_-)$, where A is the (typically huge, but frequency-dependent) open-loop gain. In this model, we neglect issues such as input offset and linearity limits.

Suppose that this device is used in the voltage-follower mode, which would ordinarily give $V_{out} = V_{in}$. Now we model the noise behavior of this amplifier. We imagine a referred-to-input voltage noise density V_n as an actual broadband white-noise emf, functionally in series with one of the amplifier inputs; and we also imagine that the very real d.c. bias current emerging from both inputs has atop it a white noise current source, with current noise density i_n. If the signal source is a resistor R, we have a circuit

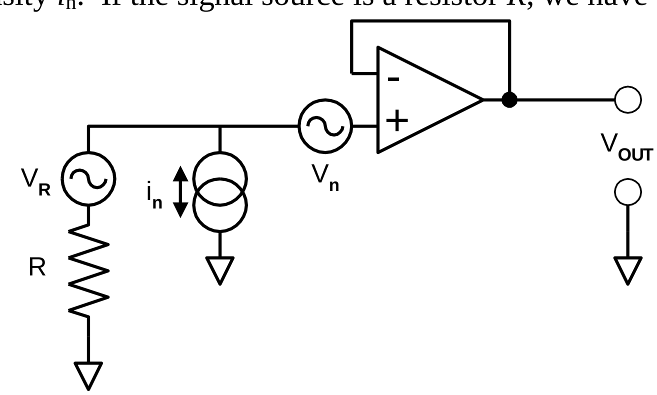

Fig. A.3b: An operational amplifier as voltage follower, showing model noise sources.

The noise behavior of this circuit includes three terms:

- amplifier voltage noise V_n is effectively applied to the non-inverting input, and (in this circuit) appears with gain (+)1 at the output.
- the Johnson noise of the resistor V_R , of noise power density 4 $k_B\,T\,R$, is also applied to the non-inverting input, and also appears with gain (+)1 at the output.

- current noise i_n , which has nowhere to go but through resistor R, where it causes a voltage drop across R which acts just like a voltage noise signal.

So the output has fluctuating voltages from three sources, presumed to be uncorrelated. As usual, the mean-square fluctuations of V_{out} simply add, to give

$$<V_{out}^2> = V_n^2\,\Delta f + 4\,k_B\,T\,R\,\Delta f + (i_n\,R)^2\,\Delta f\,.$$

Thus the noise density at the output can be written as

$$<V_{out}^2> / \Delta f = V_n^2 + 4\,k_B\,T \cdot R + i_n^2 \cdot R^2\,.$$

This is a quadratic function of R, and it is well imagined in a log-log plot vs. R. For small source resistance R, the V_n^2 term dominates; for large R, the $i_n^2\,R^2$ term dominates. These two terms make equal contributions when $V_n^2 = i_n^2\,R^2$ or at $R = V_n/i_n$. But in addition to these R^0 and R^2 terms, there is an R^1 contribution from Johnson noise, which can *exceed* the other two (amplifier-noise) terms in an intermediate-R region. If you're trying to study Johnson noise, you'd like the $4\,k_B\,T\,R$ term to dominate both the V_n^2 and $i_n^2\,R^2$ terms, at least in the neighborhood of this R-value.

To be concrete, suppose that a generic (FET-input) op-amp is characterized by input voltage-noise density V_n = 10 nV/√Hz and input current noise density i_n = 10 fA/√Hz. The quotient V_n/i_n = 10 nV/10 fA = 10^{-8} V/10^{-14} A = 10^6 Ω defines the 'sweet spot' at the crossing of the R^0 and R^2 lines in the plot. So at source resistance R = 1 MΩ, the terms V_n^2 and $i_n^2\,R^2$ both contribute 10^{-16} V^2/Hz to the noise power density. That defines the amplifier-noise baseline, against which the Johnson noise has to compete. For a 1 MΩ source resistor, that gives a density

$$4\,k_B\,T\,R \sim (1.6 \times 10^{-20}\ \mathrm{J})(10^6\ \Omega) \sim 1.6 \times 10^{-14}\ \mathrm{V^2/Hz} = 160 \times 10^{-16}\ \mathrm{V^2/Hz}.$$

Sure enough, at (and around) this source impedance, Johnson noise dominates, 160-fold in power, over both voltage noise and current noise in the amplifier.

The numbers for V_n and i_n picked above are typical for rather generic FET-input op-amps. But there are also op-amps whose front-end components are BJT-based, bipolar junction transistors. Such devices can offer smaller voltage noise (eg. 3 nV/√Hz), but they display much larger current noise (eg. 1 pA/√Hz). So a BJT-input op-amp would have its 'sweet spot' in the vicinity of a source resistance $R = V_n/i_n$ = 3 nV/1 pA = 3 x 10^3 Ω = 3 kΩ, where both terms contribute $V_n^2 = (i_n\,R)^2 = 10^{-17}$ V^2/Hz. Relative to that amplifier-noise baseline, a 3-kΩ source resistor generates Johnson noise of a spectral power density (1.6 x 10^{-20} J)(3 x 10^3 Ω) = 4.8 x 10^{-17} V^2/Hz. Again, there's a zone in which Johnson noise dominates over both forms of amplifier noise, though not by so large a factor as in the FET-based example. Then again, the absolute amplifier noise density is lower.

These examples teach us some lessons. If we have a voltage source, it'll have some characteristic source impedance. If that impedance is low (< 10^4 Ω), then a BJT-input op-amp is better suited; if that impedance is high (>10^5 Ω), then an FET-input op-amp is better suited. If (as in Johnson-noise experimental investigations) the source impedance has to vary over a wide range, then we have to understand that amplifier voltage noise, or current noise, might dominate over the source's Johnson noise in some regions of R-space.

There's another lesson to be learned. It might be that source resistance near 1 MΩ is best suited to the noise characteristics of a FET-input op-amp, but that does *not* tell us what bandwidth we can achieve. Given even 10 pF of input capacitance, a 1-MΩ source impedance gives $2\pi\, R\, C \sim 10^{-4}$ s, and a 'corner frequency' of about 10^4 Hz, 10 kHz, beyond which point the noise will roll off badly. So the Johnson noise of a 1-MΩ resistor is indeed detectable, but an experimenter might be well advised to use only the 0-1 kHz bandwidth in which to detect it.

There are finer points, too. The values V_n and i_n provided by the manufacturer are typically quoted as densities near 1 kHz. In practice, V_n tends to rise at lower frequencies (excess or $1/f$ voltage noise near d.c.). In practice, i_n tends to rise, badly for FET-input op-amps, at higher frequencies. So in addition to the 'sweet spot' of source resistance, an amplifier can have a range or region in frequency space for which it offers its lowest-noise performance. The clever experimenter (using, for example, lock-in detection) will want to ensure that the signal being investigated has been arranged to lie near the optimal location on both the source-impedance and the signal-frequency axes.

Defining 'noise temperature' and 'noise figure' of an amplifier

The noise model above also allows us to define a figure-of-merit for an amplifier called the 'noise temperature' T_n. We imagine that we have a sensing resistor R, at a temperature T_R, and we seek to detect temperature changes in T_R via Johnson-noise measurements. For our model amplifier, the noise at the amplifier's output will be the same as if we'd used an ideal (noiseless) amplifier, whose input was driven by a noise density

$$S = \langle V^2 \rangle / \Delta f = V_n^{\,2} + 4\, k_B\, T_R \cdot R + i_n^{\,2} \cdot R^2 .$$

We've seen that the Johnson-noise term dominates the amplifier-noise terms most dramatically if we pick R's value at the 'sweet spot', choosing $R = V_n/i_n$. In this case we get

$$S = V_n^{\,2} + 4\, k_B\, T_R \cdot R + i_n^{\,2} \cdot (V_n/i_n)^2 = 2\, V_n^{\,2} + 4\, k_B\, T_R\, R .$$

If we had the sense resistor at absolute zero ($T_R = 0$), we'd get the first term only; it's the net amplifier-noise contribution. Now we <u>define</u> the noise temperature of the amplifier, T_n, to be that resistor temperature at which the second (Johnson-noise) term would rise to be equal in value to the first term. So by this definition, raising the resistor temperature from 0 to T_n will raise S from $2V_n^{\,2}$ to double this value. This definition gives us

$$2\, V_n^{\,2} \equiv 4\, k_B\, T_n\, R \text{ , or } T_n = (2\, V_n^{\,2}) / [4\, k_B\, R] = (V_n^{\,2}) / [2\, k_B\, (V_n/i_n)] ,$$

so finally the amplifier noise temperature is given by

$$T_n = (V_n\, i_n) / (2\, k_B) .$$

To be concrete, we suppose that an FET-input op-amp will give us noise performance (at least in the vicinity of 1 kHz) characterized by $V_n \approx 8$ nV/√Hz and $i_n \approx 6$ fA/√Hz. Then we get

$$V_n\, i_n = (8 \times 10^{-9}\ \mathrm{V}/\sqrt{\mathrm{Hz}})\ (6 \times 10^{-15}\ \mathrm{A}/\sqrt{\mathrm{Hz}}) = 48 \times 10^{-24}\ \mathrm{W/Hz} ,$$

and we get an amplifier noise temperature of

$$T_n = (V_n\, i_n) / (2\, k_B) = (48 \text{ x } 10^{-24} \text{ W/Hz}) / (2 \cdot 1.38 \text{ x } 10^{-23} \text{ J/K}) = 1.7 \text{ K} .$$

This is remarkable performance for an amplifier whose *physical* temperature is 300 K.

It does *not* follow that a ΔT of 1.7 K is the smallest change in temperature that this resistor/amplifier combination can detect. We've defined T_n such that (compared to a resistor at $T_R = 0$), a resistor at $T = T_n$ will *double* the value of measurable noise density S. Rather than such a 100% rise in noise density $\langle V^2 \rangle / \Delta f$, it is certainly possible to detect a 10% or even a 1% increase in S. The smallest temperature change you could detect by this system would ultimately depend on

a) how stable your system would be against systematic variations, and

b) how long you were willing to wait, in averaging down the statistical fluctuations in the noise you observe.

One example of the state-of-the-art in such ΔT measurements comes from the microwave radiometry of the cosmic (blackbody) background radiation by various satellite missions. Those measurements of the 2.7-K blackbody radiation are conducted with microwave amplifiers whose noise temperatures are of order 60 K, yet they have by now resolved *micro*Kelvin variations in the blackbody temperature (variations with respect to angle, not as a function of time). But they required about a *year* of averaging time to achieve this.

Noise temperature is the preferred measure of amplifier noise performance in radio and microwave regions of the spectrum, because such amplifiers are optimized for source impedance of a fixed value (typically 50 Ω). With such an R-value matching the quotient V_n/i_n, and a noise temperature given via the product of V_n and i_n, it's clear that an assumed R-value, and a quoted noise temperature T_n, fully characterize the noise performance of the amplifier. In the world of operational amplifiers, there's no need to stick to a single source resistance, so rather than specify an amplifier by optimum source resistance and a noise temperature, the two parameters V_n and i_n are quoted instead.

In regimes where impedances *are* assumed, and noise temperatures alone therefore suffice to characterize amplifiers' noise, another figure of merit often quoted is the 'noise figure', defied by a temperature ratio, and transformed to decibel (dB) units via

$$NF = 10 \log_{10} (1 + T_n/290 \text{ K}) .$$

Our op-amp example above, with $T_n = 1.7$ K, gives $NF = 0.025$ dB, which is (very roughly speaking) a measure of how much worse is the signal-to-noise ratio at the output of such an amplifier, compared to the signal-to-noise ratio at the input.

Appendix A.4. Front-end amplifier choices and consequences

The pre-amp module in the low-level electronics part of TeachSpin's Noise Fundamentals has a 'front end', or first stage, which is user-configurable. In particular, the operational-amplifier chip for the first stage can be changed, and so can the 'topology' or choice of circuit. Here's a summary of what can be changed, and why you'd want to change it.

The choice of chip is basically between an FET- or BJT-input op-amp chip. The unit is shipped with an FET-input op-amp in place, with performance of the sort described in Appendix A.3. The voltage-noise level of the input stage is not as low as it could be made, but the range of source impedances for which this choice is adapted have led us to choose it as the default condition of the pre-amp.

If you want the lowest in amplifier voltage-noise levels, and are willing to work in the range of source impedances under about 10 or 100 kΩ, then it can help to use a BJT-input op-amp chip. The substitution is easy to make, as we've provided a pin-compatible integrated circuit in the spare-parts bin. You'll need to know how to use a 'chip puller' to removed the as-shipped input-stage chip from its socket, and you'll need to be able to recognize the pin-1 end of the 8-pin dual-inline package of the new chip to orient it properly in the socket. (Any op-amp with a '741 pinout' and tolerating ±12-V supplies may be used.)

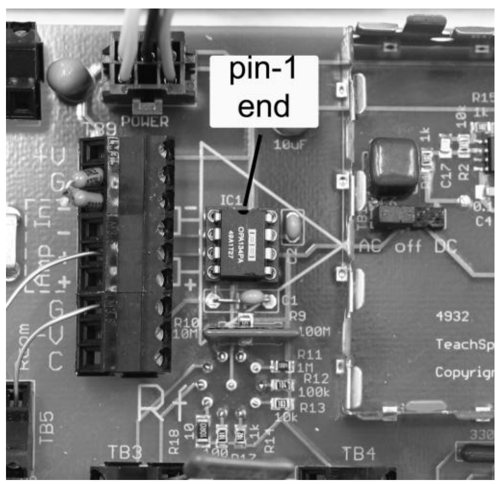

Fig. A.4a: The input-stage op-amp in the pre-amp, with the pin-1 end of chip (and socket) indicated by semi-circular 'dimples'.

You can store the op-amp chip that's <u>not</u> in use in the conductive foam in the spare-parts box.

Whether you use one chip or the other, there remains the choice of circuit topology for the first stage of the pre-amp. The unit is shipped with the configuration of a non-inverting amplifier, whose chief benefit is its very high input impedance:

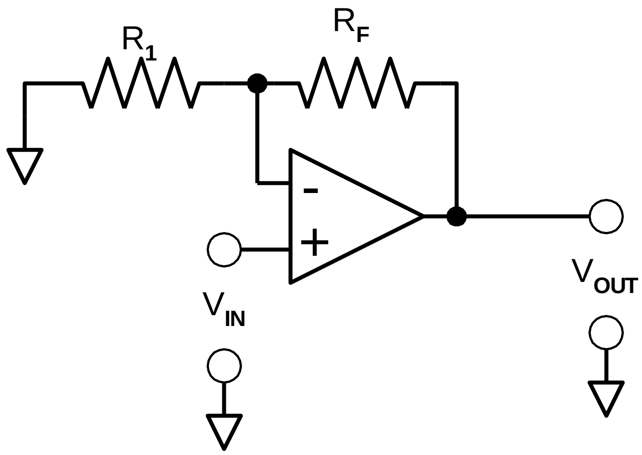

Fig. A.4b: A non-inverting amplifier topology.

This amplifier has d.c. gain $g = 1 + R_f / R_1$, which takes on the value 1 + (1.00 kΩ/200. Ω) = 6.00 in the as-shipped condition. Note that R_f is selected via the front-panel selector switch, while R_1 is a resistor attached to the terminal blocks.

A second topology retains R_f, omits R_1, and acts as a current-to-voltage converter:

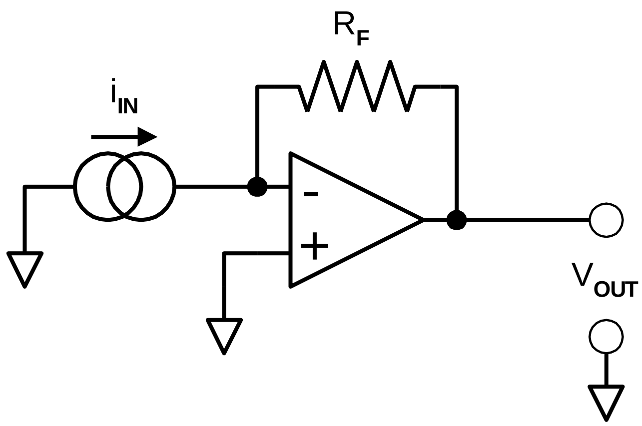

Fig. A.4c: A current-to-voltage converter topology.

A third topology is another voltage amplifier, this one inverting in character:

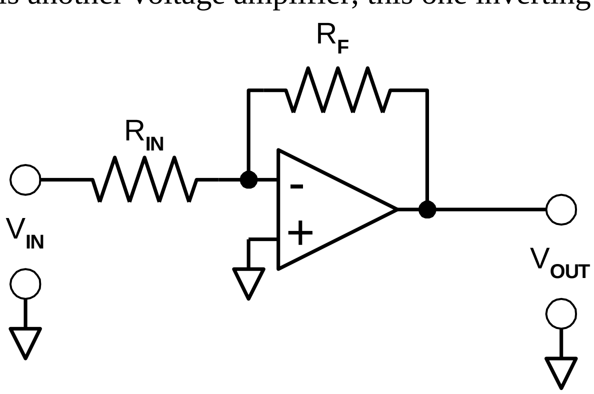

Fig. A.4d: An inverting-amplifier topology.

Here the gain is $g = -R_f / R_{in}$, and the main disadvantage is the relatively low input impedance of the circuit.

Beyond these textbook results, it is now necessary to consider the noise performance of these circuits. We take up this topic at two levels of treatment: first, the low-frequency behavior, and second, the behavior at higher frequencies (where capacitances, and op-amp bandwidth limits, start to matter).

The simpler treatment of the non-inverting voltage amplifier of Fig. A.4b is to model the op-amp voltage noise V_n as a series emf in (one of) the amplifier inputs, and to add Johnson noise as a model emf in series with each resistor. (This model omits the op-amp current noise.)

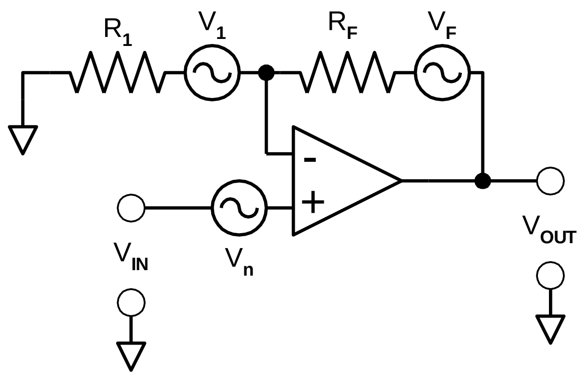

Fig. A.4e: A noise model for the non-inverting amplifier.

The result is to give

$$V_{out} = g\,(V_{in} + V_n) - V_f + V_1\,(R_f / R_1)\ ,$$

where g is still the d.c. gain given by $1 + R_f / R_1$. Relative to the expected output $g \cdot V_{in}$, the actual output displays noise of mean-square size

$$\langle \delta V_{out}^2 \rangle = g^2 \langle V_n^2 \rangle + 4\, k_B\, T\, R_f\, \Delta f + (R_f / R_1)^2\, 4\, k_B\, T\, R_1\, \Delta f\ .$$

If the amplifier noise is modeled by a voltage noise density D, or noise power density $S = D^2$, this gives output noise density

$$\langle \delta V_{out}^2 \rangle / \Delta f = g^2\, S + g \cdot 4\, k_B\, T\, R_f\ .$$

Typical values applicable to experiments in Section 1 are $g = 6$ and $R_f = 1$ kΩ; the choice of an FET-input op-amp might give $D = 8$ nV/√Hz or $S = 64 \times 10^{-18}$ V^2/Hz. Then we get

$$\begin{aligned} \langle \delta V_{out}^2 \rangle / \Delta f &= 6^2 \cdot (64 \times 10^{-18}\ \mathrm{V^2/Hz}) + 6 \cdot (1.63 \times 10^{-20}\ \mathrm{J})(10^3\ \Omega) \\ &= (2304 + 98) \times 10^{-18}\ \mathrm{V^2/Hz}\ , \end{aligned}$$

which shows that the circuit's output noise density is dominated by the op-amp's own voltage noise. The Johnson noise of the two resistors adds a small, and constant, correction -- that's why the resistors were chosen to have small values. The whole noise budget can be treated as 'amplifier noise', and subtracted by the methods of Section 1.3.

A similarly simple treatment of the i-to-V converter of Fig. A.4c is to consider the circuit with amplifier voltage noise, and resistor Johnson noise, added.

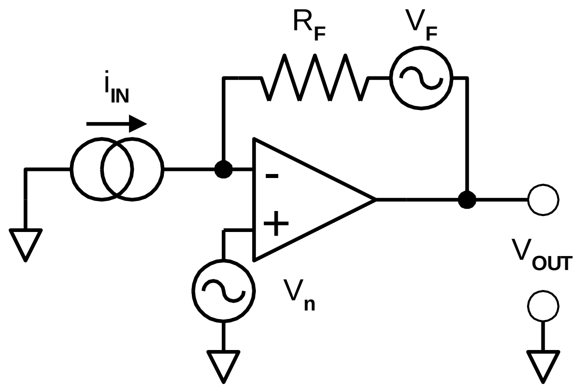

Fig. A.4f: A noise model for the current-to-voltage converter.

This model gives

$$V_{\text{out}} = - i_{\text{in}} R_{\text{f}} + V_{\text{n}} + V_{\text{f}} ,$$

and it gives output noise, relative to the expected d.c. value, of

$$<\delta V_{\text{out}}^2>/ \Delta f = \; < V_{\text{n}}^2 >/ \Delta f + < \varepsilon_{\text{f}}^2 >/ \Delta f = S \; + 4\, k_{\text{B}}\, T\, R_{\text{f}} \;\; ,$$

still ignoring the op-amp current noise. For shot-noise measurements typical of Section 3, the feedback resistor is neither fixed nor small, so we'll consider the exemplary case of $R_{\text{f}} = 10^7\ \Omega$. Then at room temperature we find $4\, k_{\text{B}}\, T\, R_{\text{f}} = (1.63 \text{ x } 10^{-20}\text{ J})(10^7\ \Omega) =$ $1.63 \text{ x } 10^{-13}\ \text{V}^2/\text{Hz} = 0.163 \text{ x } 10^{-12}\ \text{V}^2/\text{Hz}$, which dominates, by far, the op-amp voltage noise contribution of $S = (8\ \text{nV}/\sqrt{\text{Hz}})^2 = 64 \text{ x } 10^{-18}\ \text{V}^2/\text{Hz} = 0.000\ 064 \text{ x } 10^{-12}\ \text{V}^2/\text{Hz}$.

Given so large a Johnson-noise contribution from the feedback resistor, it's worth comparing its effect with the expected shot noise of the input current. A feedback resistor of $10^7\ \Omega$ is an appropriate choice for an input current in the vicinity of $i_{\text{dc}} \approx 0.5\ \mu\text{A}$, and it will give a d.c. output of $(-)i_{\text{dc}} R_{\text{f}} = (-)5$ V. Such a d.c. current allows a computation of expected shot-noise current noise $(2\, e\, i_{\text{dc}}\, \Delta f)^{1/2}$, or a current noise density $\sqrt{(2\, e\, i_{\text{dc}})} = 4 \text{ x } 10^{-13}\ \text{A}/\sqrt{\text{Hz}}$. The i-to-V converter maps this to an output voltage noise density larger by the factor R_{f}, giving $4 \text{ x } 10^{-6}\ \text{V}/\sqrt{\text{Hz}}$, or a noise power density of $16 \text{ x } 10^{-12}\ \text{V}^2/\text{Hz}$. *This exceeds, by 100-fold, the Johnson-noise contribution of the resistor, which in turn exceeds, by far, the voltage-noise contribution of the op-amp*.

That completes a 'first level' treatment of expected noise levels; at this level, resistors' Johnson noise and amplifier voltage noise have been included, but capacitance of devices, and bandwidth limits of op-amps, have not been included. We now take up some examples where these effects are considered.

We return first to the non-inverting topology of Fig. A.4b, but now include the effect of source capacitance C_{in}, in parallel with a source of impedance R_{in}.

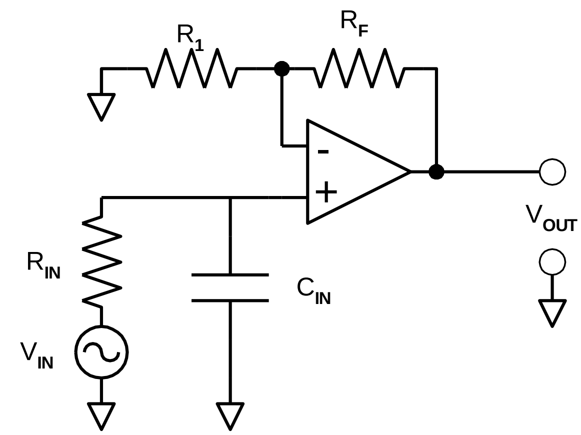

Fig. A.4g: The effects of input capacitance in the non-inverting amplifier.

We're still ignoring any capacitance that might be in parallel with R_f or R_1, because the input capacitance C_{in} is typically the effect that becomes important first. In this circuit, any Johnson (or other) emf in series with R_{in} is RC-filtered by the R_{in} C_{in} combination, which puts a bandwidth 'corner' at $f_c = (2\pi\ R_{in}\ C_{in})^{-1}$. The result is that the output noise spectrum drops *below* the white-noise limit at and above f_c, with consequences that are explored quantitatively in Appendix A.8. The input capacitance does *not* reduce the equivalent bandwidth of the op-amp voltage noise or the Johnson noise of the resistors R_1 and R_f.

A second example of the effects of capacitance is in the i-to-V converter topology of Fig. A.4c, now shown with an actual current source, having parallel capacitance C_{in}. Also shown is a user-selectable capacitance C_f in parallel with the feedback resistor R_f .

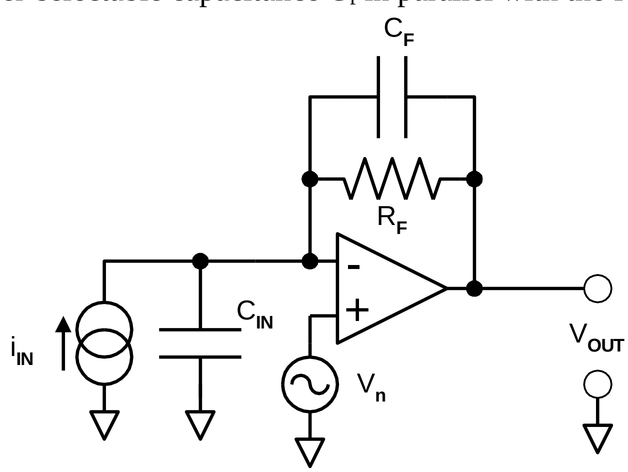

Fig. A.4h: The effects of capacitance in the current-to-voltage converter.

In the applications of Section 3, the current source may be a photodiode, and R_f is chosen to lie in the range (10^3 - 10^7) Ω, depending on the light level. Temporarily ignoring the presence of C_f, the novelty in this circuit is the frequency-dependent gain applicable to amplifier voltage noise. At low frequencies, the op-amp acts like a voltage follower for noise signal V_n, and thus gives a 'noise gain' of 1. But starting at corner frequency $f_c \approx (2\pi\ R_f\ C_{in})^{-1}$, this noise gain starts to *rise* with frequency. This gives an excess gain for amplifier noise which peaks near $f_p = (f_c\, f_m)^{1/2}$, where f_m is the open-loop unity-gain frequency of the amplifier.

To be concrete about this, we estimate $C_{in} \approx 20$ pF as the combined capacitance C_{in} of the input circuitry and the reverse-biased photodiode. If we're using an intermediate resistance value $R_f \approx 10^5\ \Omega$, then $2\pi R_f C_{in} \approx 10^{-5}$ s, and so $f_c \approx 0.1$ MHz. Given an op-amp with 'gain-bandwidth product' of $f_m \approx 10$ MHz, this tells us that there is excess voltage noise in the 0.1 ~ 10 MHz range, with about a 10-fold excess in the vicinity of $f_p \approx 1$ MHz.

This sort of 'noise peaking' is certainly visible, by using a 'scope-based FFT to look at an amplified version of V_{out}, obtained even with a photodiode in the dark.

Once such noise peaking is detected, the noise peak near f_p is readily reduced, by user selection of the feedback capacitor C_f. An approximate treatment suggests a value of C_f obeying

$$\frac{f_m}{f_p} \approx \frac{C_f + C_{in}}{C_f} \quad ,$$

which in this example gives the numbers

$$10 \text{ MHz} / 1 \text{ MHz} \approx 1 + C_{in}/ C_f, \qquad \text{or} \qquad C_{in}/ C_f \approx 9, \text{ or } C_f \approx C_{in}/9 \approx 2 \text{ pF}.$$

In practice, a slightly larger value of C_f might be used; the use of a generic C_f will give an i-to-V converter whose response drops below its low-frequency limiting value of $-\Delta V_{out}/\Delta i_{in}$ of R_f, starting at a corner frequency near $(2\pi R_f C_f)^{-1}$.

Appendix A.5. Grounding, shielding and screening, and interference

There are many sources of noise, some of them fundamental and some of them just a nuisance. Happily, electronic noise arises in obedience to Maxwell's Equations, and this provides some guidance on ways to diagnose and suppress undesired forms of noise.

1) **Grounding**

There is the complicated matter of grounding, ie. in establishing the point, assigned to have $V \equiv 0$, relative to which all potentials are measured.

In the Noise Fundamentals apparatus, local ground is exhibited by the front and back panels of the HLE, by the aluminum front panel of the LLE and the module panels installed onto it, and by the metal of the thermal probe (if that is being used). The HLE and LLE grounds are connected by both the power supply cable, and the shield of the coaxial cable, that is connecting them. Maintaining good electrical contact among all these objects is important for good electrostatic screening.

All of these grounds are connected to the third-wire ground of the a.c. power line through a 10-Ω resistor, which is in place to limit ground-loop currents that might otherwise be induced in low-resistance closed paths through which there exists a time-varying magnetic flux. The separate ground lines in the power, and signal, cables interconnecting the HLE and LLE potentially form such a ground loop. To keep such effects minimal, it is useful to keep these two cables close together, perhaps by loosely twisting them around each other.

The common $V = 0$ level is thus present at the shells of all the front-panel BNC jacks on both low- and high-level electronics. That $V = 0$ level is not changed by attaching any isolated device, such as a battery-powered multimeter, to any such BNC jack. But issues *can* arise with the use of any line-powered instrument such as an oscilloscope, whose input connectors typically have their *own* idea about what ground is, established by their own connection to the line supply.

Hence this advice: make the most sensitive noise measurements with multimeters connected to, but 'scopes disconnected from, the apparatus. Try a measurement sometime with, vs. without, a 'scope connection, while monitoring the mean-square output with a DMM, to see if ground issues are affecting your noise measurements. If they are, and if you need the 'scope connection, you can minimize this effect by plugging the power cords from Noise Fundamentals, and from your 'scope, into the same outlet fixture (and **NOT** using two outlets whose 'common ground' is established somewhere unknown or far away).

2a) **Shielding and screening**

Though these terms are not always distinguished, for this discussion we'll use shielding and screening as the names for methods of blocking the effects of external magnetic and electric fields, respectively.

Here's a way to see the effects of imperfect screening. Set up a Johnson-noise or equivalent exercise in which low-level signals are produced in the pre-amp module and sent to the high-level electronics. To favor the detection of rather low-frequency noise, use both filter sections as 1-kHz low-pass filters, and pass along the signal to the main amp. Use a gain of about 400 there, and look at the main-amp output with a 'scope. Set the 'scope to 10 ms/div, and arrange for it to trigger synchronously with your local a.c. power line. You should see 'desired noise' on the 'scope; pick a vertical-axis sensitivity which keeps the noise within range.

Now change the 'scope to the averaging mode. The noise level should drop, by about $\sqrt{N}$, where N is the number of averages you're taking. What you're looking for is residual structure, signals of fixed phase with respect to a period of 16.7 or 20.0 ms (depending on whether your power is supplied at 60 or at 50 Hz). If you don't see such interference, that's good news. But to generate some (so that you can learn to recognize it), power up a soldering gun or other transformer-containing appliance, and now hold that appliance somewhere near the pre-amp. You should now be able to view some 60- or 50-Hz interference on the 'scope. Try re-orienting and re-positioning that transformer, and testing its effects near the high-level electronics too.

Such signals as you're now seeing are due to Faraday's-Law emfs, due to rates of change of magnetic flux. The magnetic fields leaking out of the transformer core are the source, and their fields are coupling to circuit loops inside the pre-amp. Once you've seen that this can happen, you'll learn

- to imagine all such sources, and move them away if possible, especially from your pre-amp. (Remember that any line-powered instrument can be a source, too.) Take advantage of the r^{-3} drop-off of magnetic fields.
- why the pre-amp box is made of thick aluminum, and of steel. Good conductors can shield against a.c. magnetic fields by virtue of the a.c. currents induced in the shield material; good ferromagnetic materials can shield well against d.c. magnetic fields (and less well against a.c. fields.) Shielding against low-frequency a.c. magnetic fields is the hardest.
- that the LLE has the greatest sensitivity to a.c. magnetic fields, and that the size of this sensitivity will be greater if your pre-amp circuits present a larger-area loop to the magnetic field. This applies particularly to the use of the temperature probe, with its wire connection between the Temperature to the Pre-amp modules.

2b) **Screening proper**

Now that you've tested for B-field effects, almost always related to line frequencies, you're ready to think about E-field effects. If you've never observed these, it's time for you to do so. You need only a 'scope and a paper clip. Unbend the paper clip to form a single stiff wire, and poke one end of the wire into the center conductor of the BNC input of your 'scope. You've built a sort of antenna, which is resistively coupled into the 'scope, but capacitively coupled to the outside world.

Set your 'scope for 1-MΩ input impedance and for automatic triggering, and look for a signal, without averaging. The signal will grow, perhaps to large (> 20-mV) size, if you touch the paper clip. (While you're touching that lead, your body is one capacitor electrode -- where's the other?) You will likely see, amid all the other interference, some sinusoidal signals somewhere in the 25-75 kHz range. Choose the appropriate sweep speed on the 'scope, and try to trigger on these signals. The typical source of such signals are fluorescent light fixtures, computer monitors, liquid-crystal displays (including your 'scope's own display!), and lots of other devices using internal sweeps, scans, and oscillators. If your 'scope has enough bandwidth, you might also see some signals near 100 MHz, due to local broadcasts of FM signals. This might also teach you to use the reduced-bandwidth option your 'scope may offer, so long as you're doing <1-MHz noise investigations.

All of these signals are capacitively coupled, so they depend on *E*-field lines terminating on your antenna and inducing charges there. Such effects are easily 'screened' by interposing a grounded conductor to serve as an alternative, and harmless, place for those *E*-field lines to end. That's why coaxial cables have a grounded outer conductor, and why the Noise Fundamentals pre-amp and probe have grounded metal exteriors. That's why the incomplete coverage of the braided outer conductor of most coaxial cables makes them somewhat 'leaky' – signals can leak out, and interference can leak in.

To see that such things matter, here's a way to see what can 'leak through a screen'. Remove your paper clip from your 'scope, and devote the 'scope again to looking at the noise signal emerging from the main amp. (Use a configuration like that of Section 1.1.) Now here's a way partially to *defeat* the screening of your pre-amp. Remove one (of the four) screws which hold the pre-amp module into the low-level electronics. Now build an 'antenna' from a few inches of plastic-insulated wire. Strip away a cm of insulation from one end of that wire, and hold that bare-metal end. Try lowering the insulated end of the wire, down through the now-open screw hole, so its bottom end is protruding into the pre-amp's internal spaces. You should see the effect of the failure of screening, as fluctuating potentials on your fingers are now capacitively coupled into the pre-amp's circuits. Once you've seen this effect, you'll understand that screening needs to be complete to be effective. You'll understand the construction of the probe better, too.

3) **Interference**

Grounding, shielding, and screening are all defenses against interference, ie. the injection, into your desired noise signal path, of other kinds of signals generated elsewhere. Once you've seen ways to detect and defeat such undesired noise, you should get suspicious of what sources can generate it.

If your 'scope is at full bandwidth, you might see effects due to local FM stations, and then worry about nearer sources of radio-frequency noise. If these are weak enough, their high frequencies puts them out-of-band for your 0-1 MHz noise investigations. But if they're strong enough, then non-linearities can make their effects show up even in the <100kHz band. So if you can identify and turn off such sources, do so.

We've mentioned the interference in the 10 - 100 kHz band that's generated by sources like the solid-state ballasts in modern fluorescent lights. Many other devices containing switching power supplies can also generate interference in this vicinity. Typically this interference lies at, or near, one single frequency. The 'dimmers' sometimes used between the line supply and incandescent lights are another source of interference, typically at harmonics of the line frequency, but extending to very high frequencies. Your high-gain noise electronics and your 'scope are the tools for detecting such effects, but your environment is unique, and it'll take some creativity and imagination to identify all the forms of interference which might be troubling you.

There's one more source of interference that you can identify, test, and avoid: it's called 'microphonics', and it shows up as signals generated by mechanical motions of conductors due to vibrations. These motions can either cause a rate-of-change of magnetic flux, or a variation in position, changing a capacitance which maps a charge into a changing potential. Either way, you can detect microphonics by watching a 'scope view of noise while tapping suspected parts of an apparatus. When performing shot-noise experiments with the light bulb, you'll see this effect if you tap the pre-amp near that black plastic block containing the bulb. When using the temperature probe, you may also see microphonics during episodes of boiling of liquid nitrogen, especially when it fills the probe. Clearly this effect puts a premium on building rigid circuits, and then not bumping them.

Appendix A.6. Trouble-shooting

You might be familiar with trouble-shooting, which ought to be a semi-systematic method of following a signal through stages of electronics, trying to identify the place where something goes wrong. Trouble-shooting a *noise* apparatus is harder -- not only because there's no 'signal' to track, but also because it can be hard to distinguish the noise you care about from extraneous noise. So here are some suggestions to follow in case things seem not to be working as they should.

a) **Connections**

The first step in trouble-shooting is to review your interconnections. Have you plugged all your tools into the a.c. line? Do you have the right cable going to a 'scope input? Are you using the correct output from the Low-level Electronics (LLE)? Have you included all the required cables interconnecting sections of the High-Level Electronics (HLE)?

Next, check the connections you have made inside the LLE, particularly in configuring the pre-amp's first stage to your measurement needs. Pull gently on interconnecting wires and component leads to ensure they are held firmly in their terminal blocks. There are also some ground connections that you must make in configuring the pre-amp. Finally, have someone else look over your connections, to see if they match the wiring diagram, and the schematic diagram, you are trying to emulate.

b) **Power**

Start with your a.c. line cord, and look for a green LED on the cord-transformer itself. Then look for the green LED on the front left of the HLE, and another green LED on the front panel of the LLE. All should be lit when your cord is connected. If the LLE is showing an un-lit LED, suspect that you might have forgotten to turn back the internal toggle switch that's accessible when you open up the LLE and 'flip' the front panel. Check the power supplies for the operational-amplifiers inside the pre-amp, by measuring (relative to ground) the potentials at the two far ends of the terminal block at the input of the first-stage op-amp. Those points should show potentials of (±) 13-14 Volts. As a further check, use a voltmeter to check that the auxiliary ±11-V power supplies in the LLE are working -- there are monitor points on the front panel for this purpose.

c) **Signal integrity**

The modules in the HLE can be tested independently, by injecting signals from a waveform generator at an input, and looking for outputs with a 'scope. If you can spot a signal at the input of a module, and nothing emerges from that module, you have identified a problem with that module.

Recall that the filter sections give gain near 1 when you're 'in band', but can give gains <<1 when you're far outside their pass-bands. Recall that the main amplifier can have its gain set in the range 10 - 10,000. At the lower end of this range, it's easy to use a 0.5-V amplitude sine wave in, to get a 5-V amplitude sine wave out. At the high end, the gain is so large that any input amplitude > 1 mV will saturate the output.

d) **Saturation**

In normal operation, you have access to the noise signal at many points in the amplification chain: early in the pre-amp, and again at its output; and then again at every interstage connection in the HLE. You can use a 'scope, with its input set to d.c. coupling, to look for three kinds of pathologies:

> First-stage effects: you always have d.c.-coupled access to the output of the first-stage op-amp in the pre-amp. This output is connected (through a 1-kΩ resistor) to the MONITOR BNC jack on the Pre-amp panel. If this output shows a potential near (±)12 Volts, that suggests that the first stage has 'railed out', most likely because of an incorrect wiring of this first stage. Under these conditions, the first stage can*not* be faithfully transmitting noise to subsequent stages.
>
> d.c. offsets: If you use d.c. coupling between stages, then a gain stage can turn an input of (1 V of offset, + noise) into an output of (say) 10 x (1 V of offset, + noise) = 10 V of offset, + 10 x noise. Much more of this, and the d.c. offset will drive the noise into the 'rails', the upper and lower voltage limits, of about ±12 V. Once a d.c. offset has caused a signal to 'rail out', any noise atop the d.c. offset is wiped out.
>
> The third thing to look for is any evidence of noise that's gotten too large. Supposing that the use of a.c. coupling between stages (see Appendix A.2) has dealt with d.c. offsets, yet still there remains the problem of saturation. If a noise input falls in the ±2-V range, a gain-of-10 amplifier ought to produce output noise in the ±20-V range. But in this apparatus it *won't*: the largest positive and negative excursions will be 'clipped' at the levels near ±12 Volts imposed by the range of linearity of the amplifiers. That clipping not only removes part of the energy which should be in the noise, it also generates distortion, which puts energy at unpredictable locations in frequency space.

e) **Excess noise**

There will be times when you suspect that you're getting more noise than you should be. Here are some possible causes:

First, you should monitor the squarer's time-averaged output with a digital multimeter (DMM) as a measure of the noise. Then you should disconnect any and all ground-reference test instruments (like oscilloscopes) from interacting with the apparatus. If the DMM reading changes, then you can suspect some interference (typically, a ground loop) is contributing to what you're seeing. See Appendix A.5 for details.

If you have a 'scope attached, and have shown by this test that it's *not* contributing to the measured noise, you can now use that very 'scope to look for interference in what should be a pristine noise signal. Appendix A.5 teaches you the use of a.c. line triggering, plus averaging, to see the effects, if any, of all 60- (or 50-)Hz periodic effects.

You may have seen, in Appendix A.5, that a dominant form of electrostatically-coupled interference in your lab is at 25 or 48 kHz, or some other medium-high frequency, generated by fluorescent light fixtures, etc. Here's a 'scope-based way to search for contamination of your noise signal by interference from such a source:

- use the 'paper clip method' of Appendix A.5 to get a view of that interference on ch. 1 of your 'scope, and *trigger* on that interference; also pick a time base giving several cycles' view of the interference. Now put your noise signal into ch. 2 of the 'scope. Use signal averaging. If the noise signal is contaminated by this sort of interference, upon this kind of triggered signal averaging, ch. 2's signal will average *not* to zero, but to a non-zero trace revealing the interference which is contaminating the noise.
- or, use the 'scope-based FFT on the paper-clip pick-up signal to establish where (in frequency space) the interference is located; now switch to an FFT of the noise signal, and look for a peak, a location of excess noise, atop the expected white-noise background.

Turning off and on all the room lights, while monitoring with a multimeter the squarer's averaged output $<V^2>$, can sometimes show that electrostatic interference is contributing to the total noise detected. If you do see such an effect, suspect that you have a problem in some part of the LLE with imperfect screening against electrostatic effects.

f) **Suppressing interference**

If you find interference by one of these tests, you might wonder how it's getting into your system. Start by suspecting an entry point in the LLE. Check that all four thumbscrews (at top and bottom of the main LLE panel) are snugged down (finger-tight). Then check that all eight flat-head screws holding the Modules' panels to the frame are in place, and tightened. If you're not using the Thermal Probe, be sure that you 'cap off' the connector where its cable would enter. Use BNC 'shielding caps' as well at the two most crucial locations: the pre-amp's Monitor output, and the Series Resistor's Monitor position. This should deal with the potential paths for capacitively-coupled interference to get into the pre-amplifier.

Appendix A.7. Test and repair of the d.c. power supplies

The low-level electronics box, in addition to 'hosting' the pre-amp and temperature-control modules, provides you with some utilities. Among these are the two 11-Volt power supplies. (Note that one range switch, and one variable knob, control one supply with 0 to +11-Volt output, and also control another, with 0 to -11-Volt output.) These supplies have a current capability of 250 mA, yet they have been crafted to have noise levels under 5 nV/√Hz, all the way from d.c. to >1 MHz. Separate from these two utility supplies, but of very similar design, are the power supplies which run all the operational amplifiers in the pre-amp and temperature-control modules in the low-level electronics.

Building voltage supplies as 'quiet' as this takes careful regulation, which rejects voltage variation at all sorts of frequencies, and therefore has to react in a time << 1 μs. The regulators, in turn, are protected against damage in case the power supply outputs are short-circuited. But the protection cannot be instantaneous, so there are implications:

1) Please wire items to these ±11-V supplies only with the power switched OFF inside the Low-level Electronics. When you've 'flipped' the low-level panel to work on its inside, there's a toggle switch visible on the power-entry box, with a *red* LED to remind you when the power is on.

2) **Try to try not to short-circuit these power supplies !**

3) When you turn back on the power inside the LLE box, you can check that the red power-on LED comes back on, and you can check that two fault-mode red LEDs on the power-regulating printed-circuit board *don't* light up.

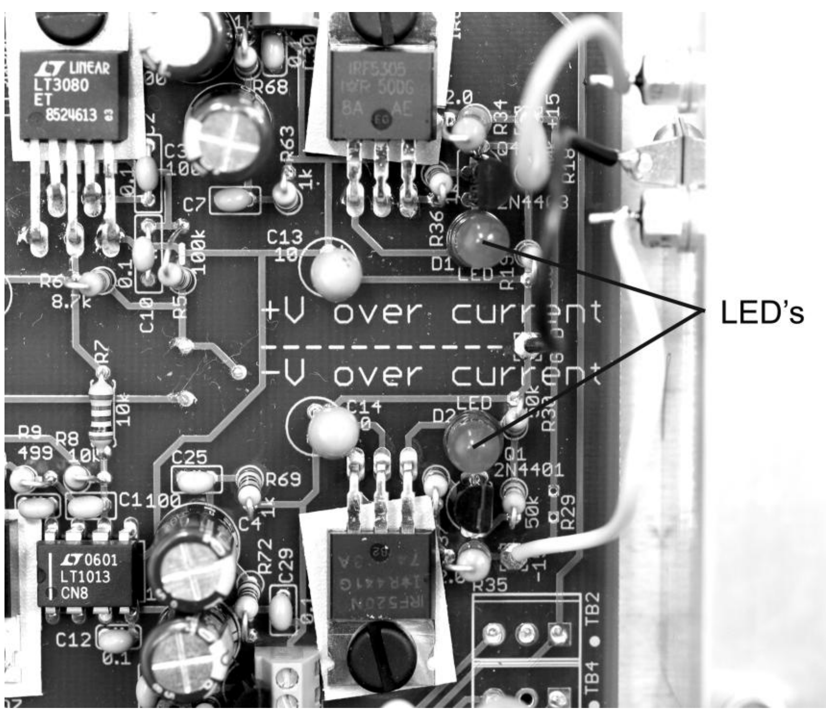

Fig. A.7a: Location of two red LEDs which are ordinarily not lit, but which will light up (though not very brightly) in case of a fault in the power-regulating circuits.

When you've flipped the panel back to its right-side-out configuration, you can re-check the proper operation of its ±11-V power supplies.

4) Some "insults" to the power supplies could result in the last-stage pass transistors failing, and failing to an open-circuit condition. Under this mode, you'll see no ±11-V capability, and you'll have to replace the pass transistors. The procedure for **Replacing** is described at the end of this s ection.

5) Other insults to the power supplies can cause the pass transistors to fail to a *short*-circuit condition. Under this mode, they will pass d.c. current, but will *fail* to remove the high-frequency fluctuations in the voltage they supply. So you'll still get an apparently useful 0 to +11-V, and/or a 0 to -11-V, output, perhaps with even a bit more voltage range than you got before. But the output you get will now be much noisier, giving voltage noise density perhaps >200 nV/√Hz instead of typically <2 nV/√Hz at the outputs. You'll need to measure this noise level to diagnose this problem, but you can't see this excess noise on a 'scope. (That's because 200 nV/√Hz of noise density, extending all the way from d.c. to 1 MHz, still gives a net voltage fluctuation of only 200 μV, rms measure.) Here's a circuit that will do the noise measurement you need -- it'll pass all noise components above about 1 Hz to the pre-amp, which is configured just as in Section 1.1:

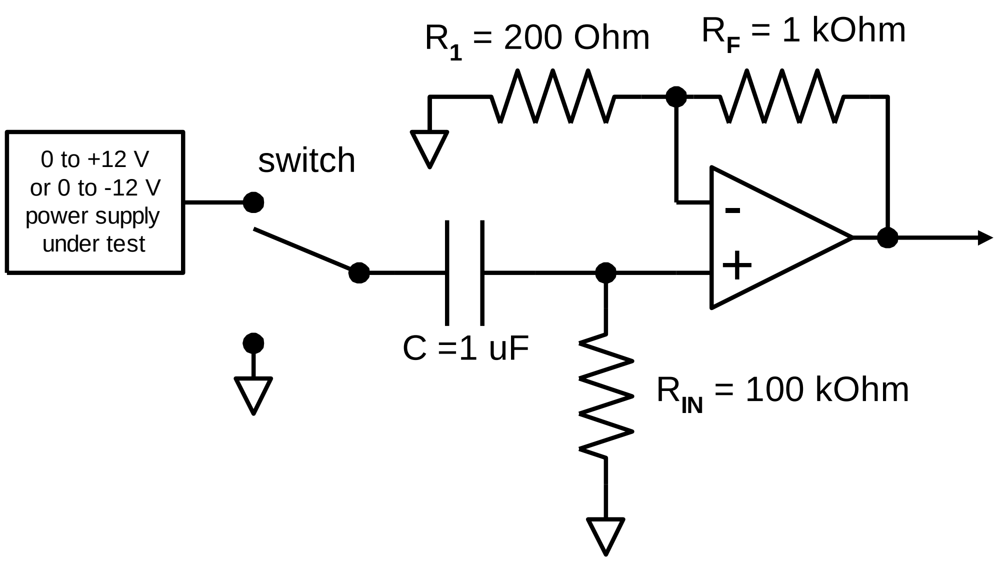

Fig. A.7b: A circuit for a.c.-coupling one of the 0-11 V power supplies to the input-stage op-amp of the pre-amp, with input stage configured for gain 6.00. (This gives overall pre-amp gain G_1 = 600.) Note that a switch setting can give you input connected to ground, rather than the positive or negative supply, as a 'control group' or check.

You'll need to understand the gain of the pre-amp, and the main-amp, and the bandwidth of your filtering, to get this method to work quantitatively. If you confirm the high-noise state of output, you'll need to replace the pass transistors.

Replacing the pass transistors

Conduct this operation with the low-level electronics' power turned **OFF**.

The transistors you might need to replace are labeled on the silk-screened printed-circuit board of the power-regulating part of the low-level electronics:

(for the 0 to +11-V supply)	Q2	type 2N4401
(for the 0 to -11-V supply)	Q6	type 2N4403
(for the op-amp + supply)	Q3	type 2N4401
(for the op-amp - supply)	Q5	type 2N4403

Before you take out a suspect transistor, make a sketch of its orientation of its black plastic package, so that you can orient the three leads of the replacement part correctly. Note that the orientations of Q5 and Q6 differ from those of Q2 and Q3. When you're ready, unscrew the three terminals to remove a transistor's leads; get a replacement device from your spare-parts bin; bend its leads to match those of the suspect device; and insert and screw it into the terminal block.

After you've replaced one transistor, here's a two-part test to see if it is working correctly: restore power to the low-level electronics, and

- use a DMM to test the d.c. level, variable or fixed, of the power supply you've repaired, for proper operation;
- then re-wire the noise-level test above to see if the power supply output has now gotten 'quiet' -- for this test, you'll have to re-flip the low-level electronics' front panel and close up the box.

Appendix A.8. Limits to the Johnson noise spectrum

We claim that Johnson noise is white, ie. that it delivers equal amounts of energy into frequency bins of equal width. But the frequency axis extends from zero to infinity, so the total energy summed over all frequencies would seem be infinite. Clearly the Johnson noise spectrum must drop to zero at some high frequency; else we'd have an 'ultraviolet catastrophe', just as in blackbody radiation.

Nyquist's derivation of Johnson noise shows that not only the disease, but also the cure, has the same form in both problems. In blackbody radiation, the electromagnetic spectral energy density $\rho(f)$ (with units of Joules of energy, per cubic meter of volume, per Hertz of bandwidth) has a frequency dependence of the form

$$\rho(f) \propto f^{3} \, [\exp(hf / k_B T) - 1]^{-1}$$

which can be written as

$$\rho(f) \propto f^{2} \frac{f}{e^{hf/k_B T} - 1} \quad .$$

The factor f^2 is appropriate to a 3-d calculation, and it turns into a factor f^0 in Nyquist's 1-dimensional calculation of electromagnetic energy in a transmission line joining two resistors. The second factor has the same origin, and the same consequences, in blackbody radiation and in Johnson noise. It's a factor which goes to a constant at low frequencies, but drops exponentially like $\exp(-hf / k_B T)$ once we have $hf >> k_B T$. That result puts quantum mechanics into our electronics problem, and it cures our ultraviolet catastrophe. It also tells us that (if nothing else were to limit the spectrum) Johnson noise can only extend out to about $f_{max} \approx k_B T / h$. (What upper frequency limit does that set, for room-temperature experiments? How about at $T = 20$ mK?) Short of this quantum cut-off, and certainly in the range relevant to tabletop electronics, we have $hf << k_B T$, and using $hf / (k_B T) << 1$ allows us to write the spectral energy density appropriate to one dimension as

$$\rho(f) \propto f^{0} \frac{f}{e^{hf/k_B T} - 1} \approx f^{0} \frac{f}{(1 + \frac{hf}{k_B T}) - 1} \approx f^{0} \frac{k_B T}{h} \quad .$$

Notice this result is linear in temperature T, and that it is also independent of frequency. So this is the origin of the overall $f^{0} T^{1}$ or 'white', but temperature-dependent, Johnson-noise spectrum.

In practice, Johnson noise nearly always drops *below* the book-value density at *much* smaller frequencies than the quantum limit mentioned above. We model a real resistor, which displays Johnson noise, as the series combination of a Johnson-noise emf and an ideal (noiseless) resistor. What we'd like to measure, with an ideal voltmeter, is the Johnson noise voltage $V_{IN}(t)$, using the circuit in Figure A.8a.

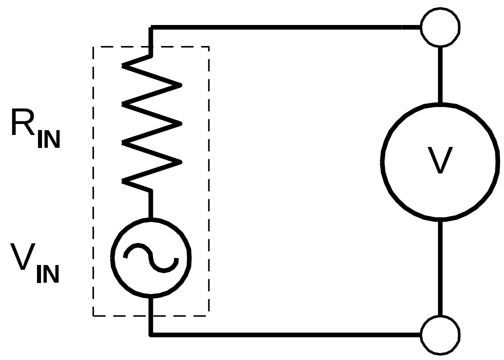

Figure A.8a

But in practice, there is capacitance, say between the two wires, or at the input of the voltmeter, so in reality the circuit we have (as shown in Fig. A.8b, in equivalent circuits) is formed by the source resistor's own resistance, and the stray (or voltmeter) capacitance. A filter like this has a 'corner frequency' given by $f_c = (2\pi\, R_{in}\, C)^{-1}$, and this corner is of real concern. If you use a source resistor of R_{in} = 100 kΩ and have even 10 pF of stray capacitance, you have $R_{in}\, C = (10^5\ \Omega)(10^{-11}\ \mathrm{F}) = 10^{-6}$ s, so $2\pi\, R_{in}\, C \approx 10^{-5}$ s, and $f_c \approx 10^{+5}$ Hz. That is to say, the Johnson noise spectrum can easily be rolling off at 100 kHz and above.

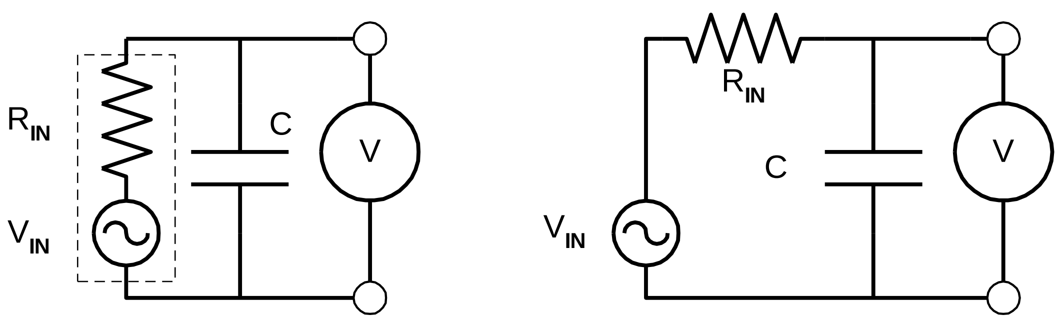

Figure A.8b

The problem is worse for larger source resistance, and much worse when the temperature probe of Section 4 is used -- there, the need for a coaxial cable raises C to about 100 pF. That's why the use of bandwidths Δf *not* extending to high frequencies is important for getting accurate values of mean-square Johnson noise.

A model for equivalent noise bandwidth, under these circumstances, is the usual integral of the square of the gain function, but now with three factors in it:

- a possible high-pass response at (low) frequency f_1 , and
- a low-pass response at (higher) frequency f_2 ; both of these taken to be ideal Butterworth functions, but now these are supplemented by
- a one-pole low-pass roll-off response at the corner frequency f_c determined by capacitive effects.

So the complete gain function becomes

$$G(f) = [\frac{(f/f_1)^2}{\sqrt{1+(f/f_1)^4}}][\frac{1}{\sqrt{1+(f/f_2)^4}}][\frac{1}{\sqrt{1+(f/f_c)^2}}] \quad .$$

You'll find (by numerical integration) that if the capacitively-caused corner f_c lies at frequency $10 \cdot f_2$ or higher, the equivalent noise bandwidth is decreased by less than 1% due to this effect.

But there is another interesting limiting case. Suppose that we use no high-pass filter at f_1, and no low-pass filter at f_2, but that the bandwidth is limited *only* by the capacitive roll-off at the source. The noise is then born with a density uniform in frequency,

$$S = \langle V^2(t) \rangle / \Delta f = 4\, k_B\, T\, R \,,$$

so the mean-square value of the emerging signal would be

$$\langle V^2(t) \rangle = \sum \frac{\langle V^2(t) \rangle}{\Delta f} \Delta f \to \int_0^\infty 4 k_B\, T\, R\, df \quad .$$

That would be infinite! but only because we left out the factor $G^2(f)$ in the integrand. In the RC-filter case at hand, $G(f)$ is the magnitude of the RC-filter's transfer function, which is given by

$$G(f) = [1 + (f/f_c)^2]^{-1/2} \,.$$

So instead of an ultraviolet catastrophe, we get

$$\langle V^2(t) \rangle = \int_0^\infty 4 k_B\, T\, R\, \frac{1}{1+(f/f_c)^2}\, df = 4 k_B\, T\, R \int_0^\infty \frac{df}{1+(f/f_c)^2} = 4 k_B\, T\, R\, \frac{\pi}{2}\, f_c \quad .$$

Now the use of $f_c = (2\pi\, R\, C)^{-1}$ gives a neatly finite result,

$$\langle V^2(t) \rangle = 4\, k_B\, T\, R\, (\pi/2)\, (2\pi\, R\, C)^{-1} = k_B\, T / C \,.$$

If you look at the circuit we're effectively using, you'll see that $V(t)$ is not only the voltage across the meter, it's also the potential difference across the capacitor C. That capacitor stores energy in amount $U_{cap} = C\, V^2/2$, so, though the instantaneous value is fluctuating, the time-averaged value of stored energy is non-zero, and given by

$$\langle U_{cap} \rangle = (1/2)\, C \langle V^2(t) \rangle = (C/2) \cdot k_B\, T / C = (1/2)\, k_B\, T \,,$$

which is a lovely illustration of the equipartition theorem. In fact, it shows that dissipation in a resistor (attached in parallel to a capacitor) comes accompanied by thermal fluctuations which prevent the resistor from discharging the capacitor all the way to zero. Instead, those fluctuations are the very mechanism responsible for the energy that, on average, is present in the capacitor.

The most remarkable feature of this result is that the measurable answer for $\langle V^2(t) \rangle$ *depends not at all* upon the value of the resistance R, yet the resistor is nevertheless the <u>source</u> of the mean-square voltage being measured. In fact, you can measure a result which depends on the resistor's temperature T, but not on its resistance R! The reason is that R's value turns up in two places, and cancels in this result: doubling R would double the Johnson noise power density, but it would also halve the equivalent bandwidth of the circuit, leading to the disappearance of R's value from the result. Perhaps you can think of a project, using a thermistor or a photoresistor to give a resistor of externally-controllable R-value, which tests this remarkable prediction.

Appendix A.9. Gaussian noise vs. white noise

You've repeatedly seen the words 'white noise', and you have perhaps also heard of Gaussian noise. Both the Johnson noise, and the shot noise, that you've been studying are ***both*** white and Gaussian in character. *But these are two separate attributes of noise*, and this section discusses the distinction.

We'll start by assuming you have a Johnson- or shot-noise source, amplified by the pre-amp, unfiltered but further amplified by the main amp to give a broadband noise signal of about 3 Volts (in rms measure).

That noise is white if it delivers equal amounts of power in any two frequency bands of equal width. That is to say: if after the main amp, an ideal sharp-edged filter were to pass all (but only) frequencies in the band ($f_0 - \Delta f/2$, $f_0 + \Delta f/2$), then the mean-square value for the resulting filtered signal would be linear in the choice of bandwidth Δf, but independent of the choice of band-center f_0. Whiteness of noise in a frequency-domain stipulation, summarized by saying *that spectral density S(f) is in fact frequency-independent*. The broadband noise you'd get in the set-up mentioned above is very near to white in the 1-100 kHz range. In practice, there might be excess low-frequency noise visible below 1 kHz, originating in the amplifiers; and there would also be some roll-off of noise at high frequencies, perhaps below 100 kHz or beyond 1 MHz (see Appendix A.8 for details).

By contrast to this frequency-domain view, the Gaussian nature of noise is specified wholly in the time domain. Think back to that broadband noise signal emerging from the main amplifier, which (in the absence of high- or low-pass filtering) has frequency content out to about 1 MHz. Suppose you sample and digitize that voltage, at a collection of random times (or equivalently, at a collection of times separated by more than the autocorrelation time of the source, which is here about 1 μs), and produce a long list of instantaneous voltage values $\{V_i\}$. Now you can make a histogram of that list, and the noise is Gaussian only if that histogram matches a Gaussian distribution. (Have you lost the rare outliers that the Gaussian distribution predicts? If so, is that because your analog voltage signal has 'run into the rails' at some point? If so, reduce the rms measure of the source noise, or the gain of the main amplifier, until even rare events will fit into your range.)

Noise can be Gaussian as a consequence of the independent operation of many independent sources; in that case, Gaussian behavior is to be expected because of the central limit theorem. Noise can be white as a consequence of processes of short or zero autocorrelation time (see Appendix A.11.) So it's no accident that some fundamental kinds of noise are both white and Gaussian.

But noise can be white and not Gaussian at all. For example, if you deliver a single pulse of fixed amplitude and brief duration, its Fourier spectrum is white (out to a frequency about equal to the reciprocal of that brief duration). Now if you have a succession of such pulses, all of identical polarity, amplitude, and still-brief duration, but occurring at random (Poisson-distributed) times, the noise this represents is still spectrally white. But its voltage histogram is *nothing* like Gaussian -- instead, it would consist of only two values, corresponding to the pulse-absent and pulse-present conditions.

Similarly, noise can be Gaussian but not white. The Noise Calibrator built into the high-level electronics has a voltage histogram which is very close to Gaussian; that's due to the central limit theorem and the use of lots of sinusoids in its construction. As is happens, the noise is also white, by design, in the 0 - 32 kHz band. But the way such pseudo-noise sources are built would allow for any desired shape of $S(f)$-function, including 'pink noise' with extra energy at low frequencies.

Finally, there's a connection between the voltage histogram of a time-domain signal and its rms measure. If $p(V)$ gives the probability of getting a particular value V for the voltage, then

$$\int p(V)\, dV = 1$$

expresses the normalization condition for probability. Similarly

$$\int V\, p(V)\, dV$$

would be the way to compute the d.c. average value of the signal (if any), and

$$\int V^2\, p(V)\, dV$$

would give the mean value of the square of the voltage. The rms measure of the signal is the square root of this,

$$V_{\mathrm{rms}} \equiv [\int V^2\, p(V)\, dV]^{1/2} .$$

Of course the rms measure is alternatively given by a calculation in the frequency domain. By definition of (single-sided) spectral density, we have

$$\left\langle V^2(t) \right\rangle = \int_0^\infty S(f)\, df \quad ,$$

so we can also write

$$V_{rms} = [\int_0^\infty S(f)\, df\ ]^{1/2} \quad .$$

The particular form of $p(V)$ for a Gaussian noise signal of rms measure A is given by

$$p(V) = \frac{1}{A\sqrt{2\pi}} \exp\left(-\frac{V^2}{2A^2}\right) \quad .$$

Thus a noise signal of 3-Volt rms measure has parameter $A = 3$ Volts, and a Gaussian distribution of voltage values which has relative size 1 (at $V = 0$), $e^{-1/2} \approx 0.607$ (at $V = 3$ Volts), $e^{-2} \approx 0.135$ (at $V = 6$ Volts), and $e^{-4.5} \approx 0.011$ (at $V = 9$ Volts). You may use an integration on the formula above to find, for example, the proportion of all voltage samples which are expected to have $|V| > 10$ Volts.

Appendix A.10. Fourier methods for quantifying noise

This section takes up the possibility of getting the frequency spectrum of noise by computational processing of an amplified noise signal, captured in the time domain. We assume an ordinary noise experiment, complete with pre-amp, filter sections(s), and main amplifier, except that in this method, we give up the use of the squarer, and instead acquire the main-amp output as a voltage-vs.-time waveform.

a) Via oscilloscopes

You've monitored the Noise Fundamentals main-amp output on an oscilloscope on many occasions, but have always viewed the waveform itself -- that's the 'time-domain' $V_A(t)$ signal. But many 'scopes offer an 'FFT' or fast-Fourier-transform utility, intended to show you a 'frequency-domain' view instead, of the spectral content of the $V_A(t)$ signal. We'll see below some details on how such things are computed from a sampled and digitized version of $V_A(t)$. Here, let's mention typical limitations of 'scope-based FFT presentations:

1) There's nothing to enforce on the user the choice of an adequate sampling rate, and the wrong choice will lead to a grossly deceptive view of the frequency content of the signal. (This effect involves the 'aliasing' of spectral content to wholly other locations in frequency space.) The requirements for sampling rates which *will* give a display faithful to the waveform's actual spectral content are given in section b) below.

   There's also nothing to prevent the user from choosing too sensitive a vertical scale on the 'scope, in which case an input signal which saturates the digitization range of the 'scope can have its spectral content splattered about unpredictably in frequency.

2) The horizontal scale of spectral displays is given correctly by oscilloscopes' FFT routines, but the vertical axis is typically left in arbitrary units. It's also traditionally plotted on a logarithmic or decibel (dB) scale, with 10 dB/div meaning that every vertical division signifying a ten-fold increase in spectral power. But reading the actual spectral power, in absolute V^2/Hz units, is a capability reserved for special 'spectrum analyzers'. One of the main goals of the sections below is to lead readers through a treatment of actual computation, by Fourier means, of results for noise power-density spectra, whose units and normalization can be understood and trusted quantitatively and in detail.

b) Sampling

This is possible given a digital sampling instrument, such as an oscilloscope, which can acquire a long series of (perhaps 10^3 to 10^5) voltage 'samples', all acquired at some uniform spacing in time. The reciprocal of this inter-sampling spacing is called the 'sampling rate', and it is critical that this rate be high enough.

How high is high enough? This comes from Shannon's 'sampling theorem', which says that *if* a waveform contains only frequency content below a maximum frequency f_{max}, *then* a sampling rate $\geq 2\, f_{max}$ is adequate. In such a case, in fact, the samples alone permit a complete reconstruction of the signal (including its unseen portions between the sampling points!). So if TeachSpin's Noise Calibrator output has frequency content (only) in the 0 - 32 kHz range, a sampling rate of $\geq$64 kSa/s (kilo Samples per second) would suffice. In practice, we might have a 'scope arranged to acquire one sample every 10 μs, giving an adequate sampling rate of 10^5 Sa/s = 100 kSa/s.

Now generic noise signals *lack* such an obvious upper-frequency limit, so for faithful sampling, it's important to limit their spectral coverage, by using a low-pass filter before the sampling. (You may have heard this called an 'anti-aliasing' filter.) But typical filters do not impose a sharp upper edge to a spectrum. You've seen in Section 2.2 that the TeachSpin low-pass filters pass some spectral energy out to $\approx 10 \cdot f_c$, where f_c is their nominal corner frequency. So if you use a 100-kHz low-pass filter, there's enough energy out to $\approx$1 MHz (and a bit more beyond) that you'd want to sample at 2 MSa/s. Note that at this sampling rate, an array of 10^5 samples will fill up in just 50 ms of time. Note also that if you use a lower corner frequency in your filter, you can afford a lower sampling rate.

c) **Scaling**

Suppose from a waveform $V(t)$ you have a collection of samples, $\{V(t_k)\}$, where the t_k are the sampling instants, separated by fixed sampling interval Δt. If there are N such samples, we could lay them out in the $-T/2 < t < T/2$ interval according to

$$t_k = -\,T/2 + (k) \cdot \Delta t \,, \quad \text{for } k = 0 \text{ to } N - 1.$$

Here T is the total duration of your sampling, and $N\,\Delta t = T$ relates N, T, and Δt.

Now if you had captured the actual continuous waveform $V(t)$, you'd reach for Fourier transforms, which we'll quote here in their complex-exponential form, and in ordinary (not angular) frequencies. In that notation, the Fourier-transform pair is

$$\tilde{V}(f) = \int_{-\infty}^{\infty} V(t)\, e^{2\pi i f t}\, dt \quad ; \quad V(t) = \int_{-\infty}^{\infty} \tilde{V}(f)\, e^{-2\pi i f t}\, df \quad ,$$

which together form a theorem, under certain conditions. But noise signals which go on indefinitely do *not* meet those conditions, since they're of constant power, rather than of finite energy. Yet we can define a scaled version of the voltage signal which preserves the frequency content of $V(t)$, via

$$W_T(t) \equiv (1/\sqrt{T})\, V(t) \,, \text{ for } -T/2 < t < T/2 \,; \qquad \text{but } \equiv 0 \text{ elsewhere.}$$

This claims that $W_T(t) = 0$ outside your sampling duration (which might be true, for all that you've recorded). With this definition, we have

$$\int_{-\infty}^{\infty} \left| W_T(t) \right|^2 dt = \int_{-T/2}^{T/2} \left| \frac{1}{\sqrt{T}} V(t) \right|^2 dt = \frac{1}{T} \int_{-T/2}^{T/2} \left| V(t) \right|^2 dt \quad .$$

Experimental voltage signals are real-valued, so this right-hand side clearly defines $\langle V^2(t)\rangle$, the mean-square value of the noise voltage, which we presume is finite. Then the left-hand side shows that $W_T(t)$ is a square-integrable function to which Fourier's Integral Theorem *does* apply, allowing us to define its transform as

$$\tilde{W}_T(f) = \int_{-\infty}^{\infty} W_T(t)\, e^{2\pi i f t}\, dt = \int_{-T/2}^{T/2} \frac{1}{\sqrt{T}} V(t)\, e^{2\pi i f t}\, dt \quad .$$

The inverse transformation is given by

$$W_T(t) = \int_{-\infty}^{\infty} \tilde{W}_T(f)\, e^{-2\pi i f t}\, df \quad .$$

Because these *W*-functions are a Fourier-Transform pair, they satisfy Parseval's Theorem,

$$\int_{-\infty}^{\infty} \left|W_T(t)\right|^2 dt = \int_{-\infty}^{\infty} \left|\tilde{W}_T(f)\right|^2 df \quad ,$$

and now we can see that *both* sides of this equation have value $\langle V^2(t)\rangle$, the mean-square noise voltage. So a physicist's viewpoint on this equality is to think of a noise source of some mean-square strength, and then to see that this given quantity (proportional to noise power) can be dis-aggregated either according to its time of occurrence (on the left), or according to its spectral distribution (on the right).

d) **Frequency content**

The Fourier transform $\tilde{W}_T(f)$ is defined on the whole line, $-\infty < f < \infty$, so it seems to contain both positive and negative frequencies. In practice, since the original signal $V(t)$ is real-valued, $\tilde{W}_T(f)$ can be shown to obey

$$\tilde{W}_T(-f) = \tilde{W}_T(+f)^* \quad ,$$

where the * stands for complex conjugation. So the information in $\tilde{W}_T$ for positive frequencies alone is sufficient to describe the whole function. It's easy to show that

$$\int_{-\infty}^{\infty} \left|\tilde{W}_T(f)\right|^2 df = \left\langle V^2(t)\right\rangle = 2\int_{0}^{\infty} \left|\tilde{W}_T(f)\right|^2 df \quad ,$$

so integrals over positive frequencies alone can tell you the full mean-square measure of the noise.

In practice, the discrete Fourier-transform methods described below are best conducted by keeping $\tilde{W}_T(f)$ as a complex function, and extracting its spectral content at the end of the computation by adding together the 'positive and negative' frequency contributions.

e) **Spectral density function**

So given a noise waveform $V(t)$, observed for a duration T, it's feasible to define a scaled function $W_T(t)$, and to compute its Fourier transform $\tilde{W}_T(f)$. Then we might define a noise power spectral density

$$S(f) = 2\,|\,\tilde{W}_T(f)\,|^2 ,$$

which is a computable function obeying the desired normalization

$$\int_{0}^{\infty} S(f)\, df = \left\langle V^2(t)\right\rangle \quad .$$

Thus the mean-square value of a voltage-noise function has been dis-aggregated into its frequency content. We'd call this $S(f)$-function the 'single-sided spectral density of noise power'. It turns out to have units of V^2/Hz, so integrating it over frequency gives Volts-squared, the correct units for $\langle V^2(t)\rangle$. So this is the computational route from $V(t)$ to a spectral density $S(f)$.

The only deficiency in this procedure is that it lacks any proper limit as $T \rightarrow \infty$. If you have an actual recording of the waveform of a noise source, and process it for ever-wider $-T/2 < t < T/2$ windows of observation, you'll find that the $S(f)$ functions computed by the above procedure gives you ever-higher spectral resolution, and shows you ever-finer details of apparent frequency variation of $S(f)$. All of this highly resolved structure is *irreproducible*, and would show up differently on a second try, for the same noise source. In practice, spot values of $S(f)$ produced by this procedure aren't convergent or useful, but wide-band or even narrow-band integrals like

$$\int_{f_1}^{f_2} S(f)\, df$$

are useful, and they *do* converge to well-behaved limits as $T \rightarrow \infty$. We'll use this fact below to motivate spectral-averaging of computed $S(f)$ values.

f) Discrete Fourier transforms

It should be clear that actual Fourier integrals cannot in fact be computed unless you were to have access to continuously-varying functions like $V(t)$. In practice, we have to be content with a finite collection of samples, such as the set $\{V(t_k)\}$ measured at N sampling points t_k separated by intervals Δt. But this very finiteness allows us to change from the integral transforms to 'discrete Fourier transform' sums instead, as follows.

We give ourselves a time window of duration T, which might be the full duration of the experiment, so that (for all we know to the contrary), a signal $V(t)$ might actually repeat, with period T, outside our window of observation. That's a convenient assumption, since any complex-valued function with period T can be written as a sum of complex exponentials of particular frequencies. We'll choose an indexing in which $f_0 = 0$ is the 'd.c.' term, $f_1 = 1/T$ is the 'fundamental' frequency, and write

$$f_n = (n)\,(1/T) = (n)\,(N\,\Delta t)^{-1}\,, \quad \text{for } n = 0, 1, 2, \ \ldots$$

Then there exists a Fourier series for the assumed-periodic $V(t)$-function,

$$V(t) = \Sigma_n\,(\text{coefficient } \#n)\,\exp(-2\pi\, i\, f_n\, t)\ .$$

Under our convenient fiction of the periodicity of $V(t)$, we can just as well take a time window $0 \leq t < T$, defining the N sampling points spaced by interval Δt as

$$t_k = (k)\,(\Delta t)\,, \quad k = 0 \text{ to } N - 1, \quad \text{still with } \ \Delta t = T/N\ .$$

To achieve a perfect fit to the N sampled data-points $\{V(t_k)\}$, it turns out that we require exactly N (complex-valued) coefficients, which we choose to write as

$$V(t_k) = \sum_{n=0}^{N-1} [\tilde{V}(f_n)]\, e^{-2\pi i f_n t_k} = \sum_{n=0}^{N-1} \tilde{V}(f_n) \exp[-2\pi i (n k)/N] \quad .$$

This sum is called the 'forward DFT', and it maps the N frequency-domain entries $\{\tilde{V}(f_n)\}$ to the N time-domain entries $\{V(t_k)\}$. This mapping is also exactly invertible -- the transformation going the 'other way' is called the 'inverse DFT', and is given by

$$\tilde{V}(f_n) = \frac{1}{N} \sum_{j=1}^{N} V(t_k) \exp[+2\pi i (k n)/N] \quad .$$

These two equations form a discrete-Fourier-transform (DFT) pair, and they are of extreme computational interest because of the amazingly efficient Cooley-Tukey 'fast Fourier transform' or FFT algorithms which have been devised to evaluate them. We've written the transforms with indices k, n = 0 to N - 1, and matched the notation and the normalization used by the open-source program Sage in its `fft()` and `inv_fft()` functions.

So here's what we actually do to get power spectral density. We want values of

$$\tilde{W}_T(f) = \int_0^T \frac{1}{\sqrt{T}} V(t)\, e^{2\pi i f t}\, dt \quad ,$$

which we compute by changing the integral to the Riemann sum we'd use to approximate it,

$$\tilde{W}_T(f) = \sum_{k=1}^{N-1} \frac{1}{\sqrt{T}} V(t_k)\, e^{2\pi i f t_k} \cdot \Delta t \quad .$$

For fictionally-periodic $V(t)$, and for a finite number of samples of it, we're content to know $\tilde{W}_T(f)$ at the frequency values f_n given above, yielding

$$\tilde{W}_T(f_n) = \sum_{k=0}^{N-1} \frac{1}{\sqrt{T}} V(t_k)\, e^{2\pi i f_n t_k}\, \Delta t = \frac{\Delta t}{\sqrt{T}} N \bullet \frac{1}{N} \sum_{k=0}^{N-1} V(t_k) \exp[2\pi i (k n)/N] \quad .$$

The factor preceding the bold dot is just $\sqrt{T}$, while the whole quantity appearing after the dot is precisely a value from an output array of N complex numbers, the result of the inverse DFT on the input array $\{V(t_j)\}$. The normalization needed after doing the DFT is just multiplication by that $\sqrt{T}$ factor, which gives $\tilde{W}_T(f_n)$-values their proper units of V · √s = V/√Hz, so we have

$$\tilde{W}_T(f_n) = \sqrt{T} \bullet \mathit{entry\ n\ of\ inv\_fft\ of\ the}\ \{V(t_k)\}\ \mathit{array} \quad .$$

Next, the absolute squares of these values give $|\tilde{W}_T(f_n)|^2$ values, with units of V^2/Hz, which are very closely related to the desired spectral density $S(f)$. We need only to remember three things:

1) Given an list of sampled voltage values, indexed by $k = 0$ to $N - 1$, namely $\{V(t_k)\}$, we need only an inverse-DFT algorithm which produces the output list, $\{\tilde{V}(f_n)\}$, a list indexed by n running from 0 to $N - 1$. The DFT algorithm needs to know the value of N (the length of the input and output lists), but it does not 'need to know' anything about the value of Δt or T. And by this stage of the computation, all reference to the W and $\tilde{W}$ functions can be dropped – the inverse DFT algorithm can be applied directly to the list of $V(t_j)$ values.

Given our choice of indexing, the frequency values associated with the index n are $f_n = (n)\cdot 1/T$, so f_0 is the d.c. term, $f_1 = 1/T$ is the 'fundamental frequency', $f_2 = 2/T$, and so on. To get single-sided spectral densities, we need to account for the 'negative frequencies' too, and (since $\tilde{W}_T(f)$ turns out to be periodic in f) these can be found in the upper half of the list of N values of f_n. In fact, to get the spectral density at a p-for-particular frequency, where integer p maps to frequency $f = p\cdot 1/T$, we take

$$S(f = p\cdot 1/T) = \left|\tilde{W}_T(f_{n=p})\right|^2 + \left|\tilde{W}_T(f_{n=N-p})\right|^2 = T\cdot[\,\left|\tilde{V}(f_{n=p})\right|^2 + \left|\tilde{V}(f_{n=N-p})\right|^2]\quad .$$

For a real-valued function $V(t)$, the DFT will produce results giving equal contributions from the two absolute-squares shown. (This is the discrete version of our previous result $S(f) = 2\,|\,\tilde{W}_T(f)\,|^2$). The frequencies which collectively account for all of the noise power include $p = 0$ (the d.c. term), and then from $p = 1$ to $(N/2)$-1. So the maximum frequency at which we get back spectral-density data is

$$f_{\max} = (\frac{N}{2} - 1)\cdot\frac{1}{T} = (\frac{N}{2} - 1)\cdot\frac{1}{N\,\Delta t} = \frac{N-2}{N}\cdot\frac{1/2}{\Delta t}\quad .$$

Since $1/\Delta t$ is the sampling frequency, we see that our spectral coverage is from d.c. to (just below) half the sampling frequency. (That's why digital audio uses a sampling frequency of 44.1 kHz, so as to cover completely the audible frequency range from d.c. to about 20 kHz.)

2) The $S(f_n)$-values thus computed obey a sum rule, which results from the discrete-Fourier-transform version of Parseval's Theorem:

$$\sum_{n=0}^{N-1}\left|\tilde{V}(f_n)\right|^2 = \frac{1}{N}\sum_{k=0}^{N-1}\left|V(t_k)\right|^2\quad .$$

This can be manipulated to give

$$\sum_{n=0}^{N/2-1}\frac{1}{T}S(f_n) = \left\langle V^2(t)\right\rangle$$

whose units match: (1/s)·V^2/Hz on the left, and V^2 on the right. This equality can be used as a valuable check on normalizations and DFT algorithms. It is also the finite-sum version of

$$\int_0^\infty S(f)\,df = \left\langle V^2(t) \right\rangle \quad ,$$

which (under the assumption of adequately dense sampling) has the Riemann-sum approximation

$$\sum\nolimits_{n=0}^{N/2-1} S(f_n)\,\Delta(f_n) = \left\langle V^2(t) \right\rangle \quad .$$

Since $f_n = (n) \cdot 1/T$, we see $\Delta f_n = 1/T$, so this result agrees with that above.

3) The $S(f)$-values thus computed will suffer from the excess spectral resolution previously mentioned, and will display a 100% scatter, with rms deviation equal to their mean. The only cure for this scatter is averaging. One way is to take M multiple successive samplings of the noise stream, each of duration T, to process each of them separately, and then to average together M multiple versions of $S(f)$. Another way is to lengthen the (single) observation window by some integer factor F, to get one long list of $S(f)$ values with high spectral resolution, and then to give up this resolution by averaging together F adjacent frequency-content readings of $S(f)$ to revert to the original spectral resolution. The extra observations, by factor M or F, will give $\sqrt{M}$ or $\sqrt{F}$ less scatter of the $S(f)$ values that result.

As an example of the latter, suppose you take $N = 2^{16} = 65{,}536$ samples of $V(t)$. (Powers of 2 are convenient, since DFT algorithms reach their highest efficiency for such array lengths.) From the methods above, you'd get back $S(f)$ values at 2^{15} distinct frequencies. If you want to end up with local $S(f)$ estimates with scatter of order 10% or less, you'll need to average together >100 $S(f)$ values. So you might average together groups of 128 high-resolution $S(f)$ values to get S-values of lower frequency resolution. But you'd still have 256 distinct $S(f)$ averages, so your frequency span would be covered with better than 1% spectral resolution.

When you finally have a table of $S(f)$ values, it is conventional to take the square root of each, converting 'power spectral density' $S(f)$ in units of V^2/Hz into 'voltage spectral density' $D(f)$ in units of V/√Hz. If your electronic signal chain has included pre-amp gain G_1 and main-amp gain G_2, then the voltage noise density at the input of the pre-amp is smaller, by factor $G_1 \cdot G_2$, than the result that you have computed via Fourier methods.

g) Confirmation

You can test the success of your computational route to spectral density by working with the Noise Calibrator signal (see Section 5.4.) You can send it, unfiltered, right into the main amplifier, set to its minimum gain, $G_2 = 10$. If you sample the amplifier's output at 100 kSa/s (so that $\Delta t = 10$ μs), you will end up with $S(f)$-values in the 0 - 50 kHz range. How many samples you can acquire depends on the storage capabilities of your 'scope, but even if you take only 10^3 samples at a time, you can make multiple 'runs' of your experiment to make the *M*-fold averaging method above work for you. Your sampling and computational route should reproduce a result close to

$$S(f) = (1.19\ \mathrm{mV}/\sqrt{\mathrm{Hz}})^2 = 1.42 \times 10^{-6}\ \mathrm{V}^2/\mathrm{Hz}$$

in the $0 < f < 32$ kHz range, and much smaller values beyond the 32-kHz limit of the noise source.

Appendix A.11. The autocorrelation function of noise

This introduces you to an alternative, and very revealing, method for viewing and thinking about noise signals. It provides a real-time method for understanding the spectrum of noise signals, or the bandwidth of circuits. We present it here first with an oscilloscope exercise you can do, and then describe its connection to the 'autocorrelation function' which is mathematically connected to the spectral distribution of noise power.

a) Observing a 'qausi-autocorrelation function'

This exercise requires only a source of noise in your Noise Fundamentals experiment, and a digital sampling oscilloscope. To do this experiment does not require the squarer, but you should set up your noise source, the pre-amp, a low-pass filter, and the main amplifier. Use a 33-kHz corner frequency for the low-pass filter, and use enough gain in the main amp to get its output up to 2~3 Volts (rms measure). Bring that signal to a 'scope, and use a vertical sensitivity that covers a ±8-V range (to accommodate the range of the noise), and use a horizontal scale of 25 μs/division.

Now if you've been seeing the noise in an automatic triggering mode, it's time to switch to a 'normal' mode, where you choose a trigger level (try a level around +6 Volts) and a slope (try positive slope, ie. trigger on signals rising through the +6-V level). You should see plenty of trigger events, since you're triggering on not-too-infrequent positive excursions of the noise. For a first look at these events, try the 'persistence' mode on your 'scope, and look for a picture like this:

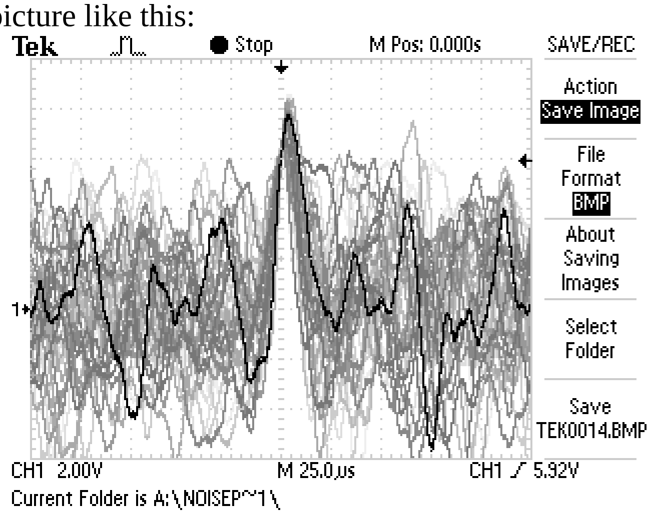

Fig. A.11a: An example of noise waveforms. Vertical scale 2 V/div, horizontal scale 25 μs/div, triggering on positive-going crossings of the +6-V level.

Notice that the trigger point has been centered on the horizontal axis. Note that every trace has the property of passing through the trigger point, both in time and in voltage. But also note that after the +6-Volt excursion, the generic trace shows signs of 'reversion toward the mean' of zero, within some finite time.

To see this in detail, change from the persistence to the 'averaging' mode of your 'scope, asking for the average of (say) 128 or 256 traces. You'll see the averaged view of reversion to the mean, with a result resembling this trace:

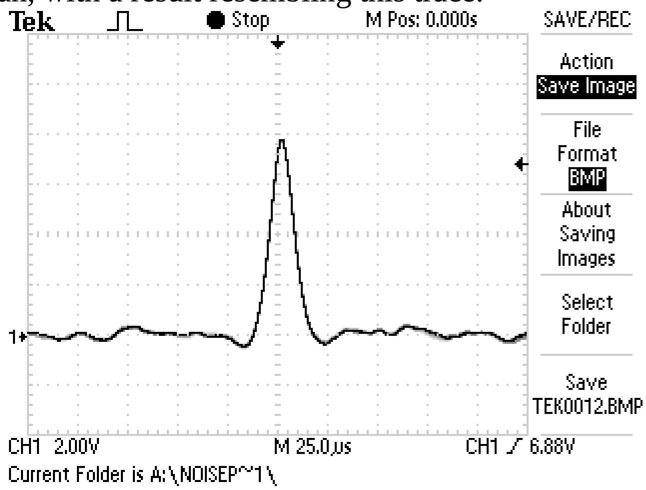

Fig. A.11b: Signal-averaged waveform in the same arrangement as above . Vertical scale 2 V/div, horizontal scale 25 μs/div, triggering on positive-going crossings of the +6-V level.

Because every individual trace passes through the trigger point, so does the average. But far enough downstream (or upstream) in time, the average value of the noise signal becomes zero again. What you have is a visual depiction of the 'autocorrelation time', which is an answer to the question 'How long does a typical positive excursion of noise last?'

To see that this pattern has something to do with your noise signal's spectral distribution, here are two comparison tests you can try:

i) change between 33-kHz and 10-kHz settings for the corner frequency of your low-pass filter. (The smaller bandwidth will give less noise power, so you may want to lower the trigger-level setting.) What you will see is a *longer* duration of the average positive excursion.

ii) change between a 33-kHz low-pass filter and a 33-kHz *band*-pass filter. (These have the same equivalent noise bandwidth, so there'll be no need to change the trigger point.) What you will see in a change in the *shape* of the result, which is due the different spectral composition of the signal you're seeing.

What are you seeing? Officially, it's called the 'conditional probability distribution', which answers the question: 'What is the ensemble-average value of $V(\tau)$, given that $V(\tau{=}0) = +6$ Volts?' [Fine point: what you're seeing is conditional on a trigger level of +6 Volts <u>and</u> a positive slope. Change to triggering on a *negative* slope to see how little difference this makes.] Happily, for Gaussian noise, this easy-to-see result on your oscilloscope is directly proportional to the autocorrelation function of the noise waveform.

b) Introducing the autocorrelation function

What lies behind the value of this oscilloscope display is its connection to the autocorrelation function or ACF called $C(\tau)$, which is defined for any signal $V(t)$. For a real-valued function, we define

$$C(\tau) = < V(t)\, V(t-\tau) >$$

where the <brackets> stand for time averaging, and where we're assuming that $V(t)$ has statistical properties which are independent of time. In this expression, τ is called the 'lag time', and this expression measures something about how different $V(t)$ and $V(t-\tau)$ can be.

It's easy to see that $C(0) = < V(t)\, V(t-0) > = < [V(t)]^2 >$ gives the mean-square measure of the noise; this shows that $C(0)$ is always positive, and also shows that $C(\tau)$'s units are Volts-squared. There are also some properties of the official autocorrelation function $C(\tau)$ which are similar to those of the 'quasi-ACF' which you're viewing on the 'scope. The first of them (established via the Cauchy-Schwarz inequality) is that $|C(\tau)| \leq C(0)$ for any choice of τ. It's also feasible to show that $C(\tau) = C(-\tau)$, which shows that $C(\tau)$ is a function symmetrical about $\tau=0$ (where it thus has an absolute maximum).

The definition of $C(\tau)$ also makes it clear why the function drops off with time, and why it distinguishes the regimes of short, vs. long, compared to some autocorrelation timescale. If τ is short enough, $V(t)$ and $V(t-\tau)$ will be similar, hence much more likely to be of the same (as opposed to opposite) signs. So the product $V(t)\, V(t-\tau)$ will be more probably positive than negative, so its time-average will be positive. By contrast, is τ is long enough, the present value $V(t)$ will be *uncorrelated* with its value 'τ ago' in the past, $V(t-\tau)$. So at those times when $V(t)$ is positive, $V(t-\tau)$ will be as likely to be negative as positive. Hence the product $V(t)\, V(t-\tau)$ will also be as likely to be negative as positive, and so its time average will be near zero.

So the shape of the function $C(\tau)$, like the quasi-ACF you saw on your 'scope, tells you about the degree to which the signal has some 'staying power' or even 'memory'. That is in general not the memory (if any) in the original source of the noise, but rather due jointly to the noise source and the bandwidth that might have been imposed upon its signal by subsequent filtering. In fact, there's an inverse relation between bandwidth and autocorrelation time: a truly white-noise signal with bandwidth to $f = \infty$ would have *zero* autocorrelation time, and act like a system with no memory at all. But the smaller the bandwidth, the longer the autocorrelation time; by the time you get to a pure sinusoid or any other periodic signal, the ACF shows non-zero correlations at arbitrarily long lag times.

And there's more than this informal connection between spectral distribution and autocorrelation function. It turns out that $C(\tau)$ on the one hand, and the power spectral density $S(f)$ on the other hand, are closely related as a Fourier transform pair. So knowing either function of this pair fully determines the other. For example, if we had a sharp-edged spectral distribution of noise, with $S(f)$ a constant from d.c. up to some maximum frequency f_m (but zero beyond that point), then taking the Fourier transform of $S(f)$ allows us to predict

$$C(\tau) \propto \sin(2\pi f_m \tau) / (2\pi f_m \tau) ,$$

which has the sinc-function 'wiggles', and displays its first zero-crossings at lags $\tau = \pm 1/(2 f_m)$. The 'wiggles' disappear for spectral distributions lacking as sharp a cutoff as in this example, but the width (which was $1/f_m$, between innermost zeroes, for this sharp-edged distribution) will continue to be inversely proportional to the bandwidth of the spectral distribution.

It also follows that two sources with distinct spectral distributions (such as the low-pass, vs. the band-pass, filtered versions of white noise) must have distinguishable $C(\tau)$ functions. It also follows that careful measurement of $C(\tau)$'s values can provide a quantitatively reliable way to compute $S(f)$. The place to pursue this connection is a presentation of the Wiener-Khinchin theorem in signal processing.

c) Best use of a 'scope's FFT-capability

If you have used the FFT utility of your 'scope to view the frequency spectrum of a noise signal, you may have been horrified at the fluctuations of the spectrum. A white-noise spectrum should give an $S(f)$-function which is a flat line, but in practice you've seen a host of jagged spikes and dips downward from the level you've expected. Appendix A.10 deals with some of the cures to this problem that you can impose if you compute your own FFTs off-line, but here's a capability which you can execute on your 'scope directly.

Ideally, you could ask for the FFT, and then the 'averaging' mode. What you'd *want* is an average of many successive spectral distributions. But what you'll *get* is the Fourier transform of (a bunch of noise waveforms all averaged together). That doesn't work right -- averaging the noise together (first) tends to wash away its strength, so your signal and the consequent FFT disappears.

So here's what to do instead. You set a trigger level as above, and average the time waveforms that appear at this trigger level, to get what we've called the quasi-ACF. Now you are using the averaging mode of your 'scope in a way which does not average the time-domain signal away toward zero; instead, you're getting a version of the auto-correlation function $C(\tau)$. *Then* you ask for the FFT, and you'll get the FFT of $C(\tau)$, and what you'll get from the computation is closely related to $S(f)$ -- by the Wiener-Khinchin theorem. What will be *displayed* may be $| S(f) |^2$, so it's hard to use this display with quantitative certainty about the scale of its vertical axis. (The half-power points might be depicted at -6-dB down, for example.) But you *will* get a display of spectral content with markedly smaller scatter, vertically, than you'd get in the direct FFT of the input waveform.

d) Using the quasi-ACF for analysis

For present purposes, here's another result which is easy to prove about $C(\tau)$, and also easy to observe using the quasi-ACF method on your 'scope. Suppose that a signal $V(t)$ is really the superposition of signals from two distinct sources,

$$V(t) = V_a(t) + V_b(t) \; .$$

You could call one of these signal, and the other noise; or it might be that one is the desired noise, while the other is *un*desired noise. Here's the $C(\tau)$ you get in this case:

$$\begin{aligned} C(\tau) &= \langle V(t)\, V(t-\tau) \rangle \\ &= \langle\, [V_a(t) + V_b(t)]\, [V_a(t-\tau) + V_b(t-\tau)]\, \rangle \\ &= \langle\, V_a(t)\, V_a(t-\tau)\, \rangle + \text{two cross terms} + \langle\, V_b(t)\, V_b(t-\tau)\, \rangle \; . \end{aligned}$$

The cross terms include $\langle\, V_a(t)\, V_b(t-\tau)\, \rangle$, which is *zero* for any and all τ-values, provided only that 'a' and 'b' stand for physically separate, ie. uncorrelated, sources of noise. So in this case,

$$C(\tau) = C_a(\tau) + C_b(\tau) \; ,$$

which shows that the ACF you'd observe is simply the sum of the ACFs you'd observe from the two sources separately.

You can get a great view of this process with your 'scope-based quasi-ACF. Try grounding the input of your low-pass filter section, set it to a 33-kHz corner frequency, and send its output to the main amp, set for maximum gain of 10^4. Observe the raw noise signal at the output of the main amp, and you'll see on a 'scope an entirely uninformative noise waveform, with no hint that it's composed of two kinds of noise. There's noise generated in the filter (with spectrum rolling off at around 33 kHz), and there's noise generated in the main amp itself (whose bandwidth extends to ≈1.5 MHz). Use your view of the net noise to choose the right vertical sensitivity, and to choose a good trigger level, and averaging, to produce a quasi-ACF. Now use sweep speed about 5 μs/div and look at the result, which should resemble this:

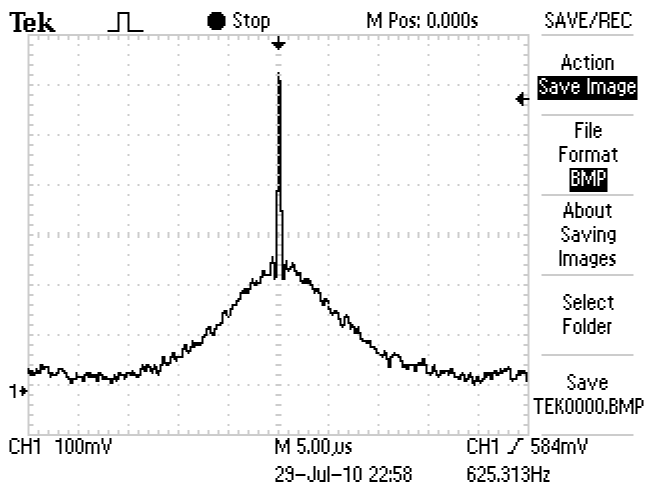

Fig. A.11c: The 'quasi-autocorrelation function' revealing the presence of two kinds of noise in the filter-plus-main-amp combination discussed above.

You can view this with a variety of timescale settings on your scope, to get a good view of both the narrow peak, and the broad hill atop which it's standing. You can even see some structure in that narrow peak -- look for some little valleys on either side of the narrow tall peak. The value of this exercise is the mental separation it permits. The narrow peak is the part of $C(\tau)$ due to a signal of short autocorrelation time, which must be of large bandwidth -- we identify that as main-amplifier noise. The full width at the base of the sharp peak is about 0.6 μs, which corresponds to a noise spectrum extending up to $f_m \approx 1.7$ MHz. The broad hill underlying the peak is the part of $C(\tau)$ due to a signal of *long* autocorrelation time; that must be of small bandwidth, and we claim it's due to noise born in the filter section. In fact the width of the broad hill is of order 20 μs, which is consistent with a frequency spectrum extending to about 50 kHz. And in agreement with the bit of theory above, we can now understand why these two pieces of the autocorrelation function simply add up to give the shape observed.

To test these claims further, you can change the bandwidth chosen on the filter -- what effect should that have? Or, you can send some white noise into the filter's input, which does not change the amount of noise that's actually generated within the main amp -- how will this show up? Or, you could imagine some interference (see Appendix A.5) from fluorescent-light ballasts, of frequency perhaps 25 or 48 kHz and approximately sinusoidal in character, is underlying your noise -- can you compute what *third* contribution to $C(\tau)$ that would create? As a practitioner, you can gain some instinctive knowledge from the easily-acquired, real-time quasi-ACF display on your 'scope, and use it as a key to diagnosing many kinds of experimental pathologies.

Appendix A.12. Fluctuations in measured noise: The Dicke limit

Noise signals are random, and as a result, measurements of 'noise power' display statistical fluctuations. This Appendix explains some nomenclature for these fluctuations, and describes and justifies the expected size of the fluctuations.

Let's imagine Johnson noise, or shot noise, measured by a now-familiar arrangements. We have an original noise voltage $V_n(t)$, characterized by zero mean but a non-zero mean-square. We pre-amplify it (by gain G_1), we filter it to bandwidth Δf (using filter gain function $G(f)$), we further amplify it (with gain G_2), and thus form a filtered and amplified noise voltage V_{in}(t) as input to a squaring circuit. The mean of $V_{in}(t)$ is still zero.

When we square $V_{in}(t)$ to get $V_{out}(t) = [\ V_{in}(t)\ ]^2 / (10\ V)$, we finally get a signal whose mean is *not* zero. So when we average it over averaging time τ, we get a non-zero average

$$V_{meter} = < V_{out}(t) > \propto < [V_n(t)\ ]^2 > \ ,$$

and we can call that concrete meter reading a 'measure of the noise power'. But we can also easily see the visible fluctuations in the meter position, which we can call 'fluctuations in the measured noise power'. (They are sometimes called the 'noise in the noise', or 'second noise'.)

How big do we expect those fluctuations to be? This question was first addressed by Dicke, in an appendix to the paper [Rev. Sci. Instrum. **17**, 268 (1946)] which introduced lock-in detection, and which used it to measure room-temperature blackbody radiation in the microwave region of the spectrum. The result is therefore called the 'Dicke radiometer limit', usually expressed as a characteristic fluctuation δT observed in an instrument whose output gives T, a radiometrically-measured source temperature. Dicke's result can be written in terms of the bandwidth Δf and the averaging time τ as

$$\delta T / T \approx (\Delta f\ \tau)^{-1/2}\ .$$

This result applies more generally than just to temperature measurement by radiometry, and it also applies to noise-power measurements as conducted in Noise Fundamentals. If $V_{meter}(t)$ is the instantaneous voltage applied to the meter, which is traceably connected to the mean-square noise signal $< V_n^2 >$ at the source, then fluctuations in the meter output are also given by

$$\delta V_{meter} / < V_{meter} > = const \cdot (\Delta f\ \tau)^{-1/2}\ .$$

Here δV_{meter} can be taken to be the standard deviation of a sample of (independent) readings of the meter. The equivalent noise bandwidth used in the filtering chain provides the factor Δf, and the averaging time used between the squarer and the meter provides the factor τ. Finally, the constant is of order 1; its numerical value depends on just what kind of time-averaging is used. (The TeachSpin equipment uses two successive one-pole filters, each of time constant τ, and the predicted value of the constant is about one-half.)

So if we measure noise using coverage limited by a low-pass filter of corner frequency 100 kHz, we have $\Delta f \approx 114$ kHz. If we choose a $\tau = 0.1$-s averaging time at the meter, we get

$$\delta V_{\text{meter}} / < V_{\text{meter}} > = const \cdot (114 \text{ x } 10^3 \text{ /s} \cdot 0.1 \text{ s})^{-1/2} = const \cdot 0.009 \ ,$$

so we expect fluctuations of order 0.9% in the meter reading. Of course we'd have to make $V_{\text{meter}}(t)$ readings at a time spacing of $\geq \tau$, or at least 0.1 s apart in time, for them to represent statistically-independent readings, in computing the fluctuation level. δV_{meter} as a standard deviation.

The dependence of this Dicke limit on Δf and τ is easily visible. Keeping τ fixed at 0.1 s, to give a set of readings with which the analog meter can 'keep up', you can try out he effect of changing from 100 kHz, to 10 kHz, to 1 kHz for the corner frequency of the low-pass filter in the high-level electronics. (Of course, when you reduce the bandwidth, you'll want to raise the gain G_2 to keep the average meter reading near 1 Volt.) What you should see is visibly *larger* fluctuations in the meter's position, since the Dicke equation predicts fluctuations, about the 1-Volt average, of order 0.9%, growing to 3% and then 9% as you reduce the bandwidth. So for the smallest *statistical* fluctuations in any noise measurement, it's always best to use the largest possible bandwidth Δf. (Of course, there may be growing *systematic* errors, such as the effects of capacitive roll-off, which accompany such a choice of larger Δf.)

A simple explanation of the reason for a Dicke limit also explains the $\tau^{-1/2}$ dependence. We know that the output of the amplifier/filter chain is limited to bandwidth Δf. It follows that the autocorrelation time of this signal is of order $1/\Delta f$. Thus the use of 100-kHz bandwidth gives a filtered noise signal with an autocorrelation time of about 10 μs. Hence there's a 'fresh value', or a statistically-independent measure of $<V_{\text{in}}(t)^2>$, available every 10 μs. If we use a $\tau = 0.1$ s averaging time, the number of statistically-independent measures of noise power we can make during that time is

$$N \approx (0.1 \text{ s of time}) / (10 \ \mu\text{s per fresh measurement}) = 10^4 \ .$$

The mean of all these 10^4 measurements is what the meter reveals via its average reading. But since those 10^4 individual readings are each of them random and independent, we expect fractional fluctuations of the meter reading to be of order $N^{-1/2}$. This hand-waving argument in fact gives

$$\delta V_{\text{meter}} / < V_{\text{meter}} > = 1 \ N^{-1/2} = 1 \ \{ \tau \ / (\Delta f)^{-1} \}^{-1/2} = 1 \ (\Delta f \, \tau)^{-1/2} \ ,$$

just as discussed above.

In a computer data-logging environment, it's easy to get a large sample of meter-reading voltages. For any choice of τ, it's easy to test if successive readings, taken at time spacing τ (or better, 2τ), display statistical independence (ie. absence of correlation). It's also easy to compute $<V_{meter}>$, and to form the histogram of V_{meter} readings (expected to be distributed about their mean in Gaussian fashion). The standard deviation of all the readings displayed in the histogram defines the characteristic scale of fluctuations, δV_{meter}. Then the Dicke limit can be tested empirically for its Δf and τ dependence.

The Dicke limit also imposes stiff requirements on any noise-based experiment that seeks to attain really high precision. If shot noise were to be used in search of a part-per-million measurement of *e*, and if all systematic effects were fully under control, this limit would ultimately require some meter reading to display

$$\delta V_{meter} / < V_{meter} > = 10^{-6} \ ,$$

and that, in turn, would require $(\Delta f\, \tau)^{-1/2} = 10^{-6}$, or $(\Delta f\, \tau) = 10^{+12}$. For a bandwidth of $\Delta f \approx 100$ kHz = 10^5 /s, that would require a total averaging time of $\tau = 10^7$ s, or about four months! (One method for doing this would be to set the meter-averaging switch to a 1-second time constant, and take one reading every second until 10^7 readings had been collected and averaged.) This provides another example of the desirability, at least on statistical grounds, of using the largest possible bandwidth Δf.