Noise Fundamentals - Measuring Johnson noise

MB 008

Preamble -Definition, kinds and uses of noise

Ordinarily speaking, noise is an unwanted sound but we can extend the definition of ‘noise’ beyond acoustics to the general field of information. Since almost any signal that is a function of time can be translated into a voltage, it is convenient to use the concept of a voltage signal. A ‘noisy signal’ is one that in addition to the expected voltage has an 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, for example a mobile phone interfering with the circuits in a radio. The kind of interference you 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 any 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, 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 in this experiment is ‘fundamental noise’, noise that is intrinsic and inevitable because of the physical nature of the apparatus. You will be observing Johnson noise, which arises from the Second Law of Thermodynamics. Noise sources like Johnson noise 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 for example Boltzmann’s constant, k_{B}, can be determined from the voltage or Johnson noise of resistors.

However, measurement of ‘fundamental noise’ has its experimental challenges. There is a saying about noise measurements: ‘you are either measuring too much or too little signal’. In this practical you are going to investigate what “just the right amount” of noise is.

1. Johnson noise at room temperature

1.0 The reasons for Johnson noise, and its predicted size

It is well known that $V = I\,R$, 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' $V_{\mathrm{J}}(t)$ across its terminals, a potential difference that has all the character of an internal (a.c.) emf built into the resistor. The electro motive force (emf) which the resistor generates is called 'Johnson noise', and it arises because of the deep thermodynamic connection between dissipation (which any resistor necessarily 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 $V_{\mathrm{J}}(t)$ is, on average, zero. But $V_{\mathrm{J}}(t)$ exhibits fluctuations, positive and negative, about that average value of zero. To quantify these, we form the (always-positive) square of $V_{\mathrm{J}}(t)$, and time-average that, giving a 'mean square' voltage which we denote as $< V_{\mathrm{J}}^2(t) >$. The predicted value for $< V_{\mathrm{J}}^2(t) >$ was first deduced by Nyquist, following Johnson's empirical discovery of the noise, and it's given by the expression

$$< V_{\mathrm{J}}^2(t) > = 4k_{\mathrm{B}}\, RT\Delta f.$$

Here k_{B} is Boltzmann's constant, T is the (absolute) temperature of the resistor, and Δf 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 root-mean-square (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.

Using Nyquist's Theorem, how large a Johnson-noise voltage should we expect from a typical resistor? If you simply hooked up a 100kΩ room temperature resistor 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 around zero, with typical excursions of order $\pm 10\mu V$. (What sort of time scale would you expect for fluctuations of this signal?) 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.

Here's a 'thought experiment' to help you see that some sort of Johnson noise should 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 a lot 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 50Ω resistor at room temperature, connected to nothing but a lossless coaxial cable of 50Ω impedance; and imagine there is another 50Ω resistor, but down in a Dewar at T = 4K, connected to the far end of this cable. If the coaxial cable is superconducting with zero thermal conductivity what happens to the temperature of the two resistors in the presence of only electrical conductivity. The 'Johnson emf' in each resistor 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

It is possible to 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. This is exactly what you will now do.

Connect power to the High Level Electronics (HLE) box, which also provides power for the Low Level Electronics (LLE) box.

Select a ‘source resistor’ of $R_{\mathrm{in}} = 100\ \mathrm{k\Omega}$ in the pre-amplifier module of 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 (see Fig 1) is wired to give a ‘gain’, or amplification factor that depends on the feedback resistor (R_{f}). For a setting of say $1\mathrm{k\Omega}$ what gain would you expect? (The feedback capacitance C_{f} is not connected in the default configuration, so its setting is irrelevant.) The graphics on the panel of the pre-amp shows that there is an additional amplification stage, with gain 100, following this first stage, what is the combined gain (G_1) for the preamp module? Now you can connect the pre-amp’s output with a coaxial cable, to an oscilloscope, to see any signal present. Use an appropriate vertical scale and sweep speed of around $5\mu s$/div and a scope trigger near zero volts.

Can you calculate/estimate approximate $V_{\mathrm{J}}(t)$ values? Why do you need the pre-amplifier?

The wiring diagram for the required configuration is shown in Fig 2. The connections indicated in thick grey lines are wires that can be reconfigured. (By contrast, connections shown in thin solid lines are already established for you on the printed- circuit boards.)

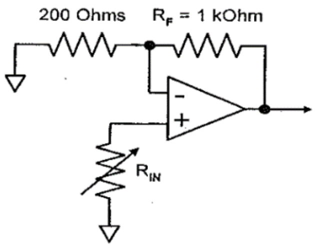

Figure 1: Johnson noise preamplifier schematic

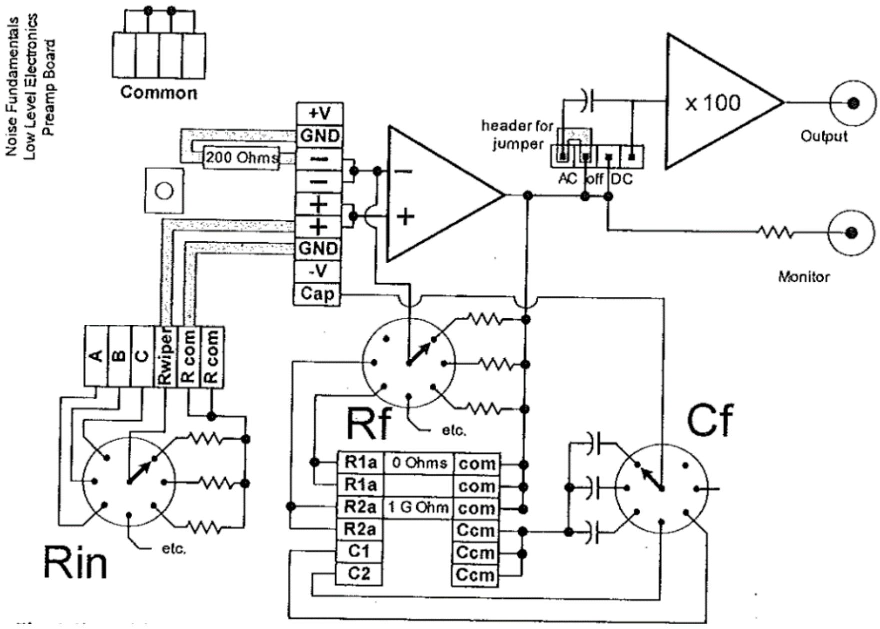

Figure 2: Wiring diagram of the default configuration of the interior of the low-level electronics (LLE).

The signals you see emerging from the pre-amp are rather small. So next use a cable to connect the pre-amp output to the HLE box instead, where you can filter and amplify the still-small noise signals. Set up the arrangement show in Fig. 3, with a frequency band, extending from about 100 Hz to about 100 kHz and AC coupling. The first filter is used as a high-pass and the second as a low-pass filter. After the output of the two filters, you have Johnson noise, pre-amplified by a known factor, and then filtered to pass only the 0.1k -100kHz frequency band.

There is another stage of "main" amplification which you should set to give a gain of around 300. Finally, at the output of this main amplifier, you'll have a signal large enough to see easily on a scope. This is Johnson noise, due to thermal driven current fluctuations.

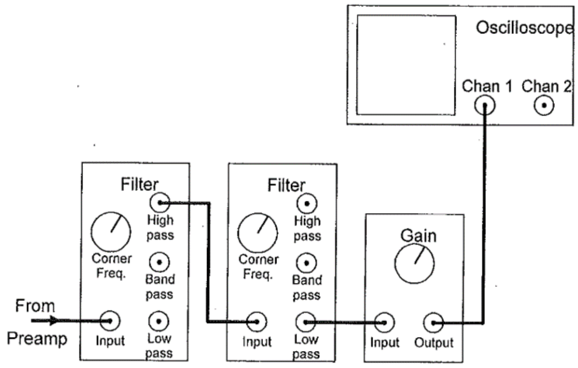

Figure 3: Cabling diagram for the high-level electronics. (left) Filter: selector to .1k, (right) Filter: selector to l00k, (a.c. coupling for both filters). Gain Fine Adjust 30, toggle x1, toggle x10

Look at a few different Δf bandwidths and R_{in} resistances. You might like to save a few traces/screen shots. *Note: If you are using Excel (or similar programs), they are designed for spreadsheets and may slow down considerably (crash) if many large data files are imported for viewing.*

To help convince yourself 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 try a few different R_{in} resistances. When you change R_{in} by a factor 10 or 100 how much does the signal change? Follow the amplification chain and satisfy yourself that the size of the signal that comes out makes sense.

In general measuring the intrinsic noise can be a really useful test when setting up sensitive measurement circuity. The background noise floor is **always** present and should match the expected value otherwise (a) something is wrong or (b) you have discovered new physics. *Usually it means (a) and you can try fixing it before looking at your actual signal and getting on with (b)*.

1.2 Quantifying Johnson noise

You have hopefully just observed a rapidly-fluctuating signal on the oscilloscope which seems to be consistent with Johnson noise but you should now quantify it. The method described here executes quite directly, in analogue electronics, the very operation built into the mean-square definition of noise.

You will be using the filters, amplifier, multiplier and output modules of the HLE box to implement this "mean-square" operation. For a moment ignore the two filter sections, (you may well be able to guess what they do) and instead connect a test signal from the function generator to the input of the Gain module of the HLE box. Try an input 1kHz sinewave, amplitude 0.5V and a module gain of 10, what output do you get? (Compare the input and output on Ch1 and Ch2 of the scope.) Now vary the input amplitude and gain and note what you observe. Briefly describe (one or two sentences or a sketch) what the gain module does and if there are any (frequency or amplitude) limits to its operation. Connect a "suitable" test input signal from the function generator to the multiplier module. What does it do (use AxA)? Are there any limits to the input signal it can deal with? What does it produce with 0, 1 & 2V input? Next look at the output module, this has an averaging effect. Use a square wave input signal (applied via the multiplier) and look at the effect of the time constant settings. What do the time constant numbers on the switch correspond to?

Now back to quantifying John son noise. Connect up the equipment as shown in Fig 4. You can turn off the signal generator for now.

As you have discovered at the MONITOR BNC (and also internally to the output module) the multiplier circuit delivers, a real-time output voltage

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

which still has dimensions Volts (due to the fixed 'scale factor' of 10 Volts in the denominator). With the preamp module and R_{in} as they source look at $V_{\mathrm{out}}(t)$ on your scope, and notice that it is always positive, unlike your input noise signal $V_{\mathrm{in}}(t)$, which has a mean of zero.

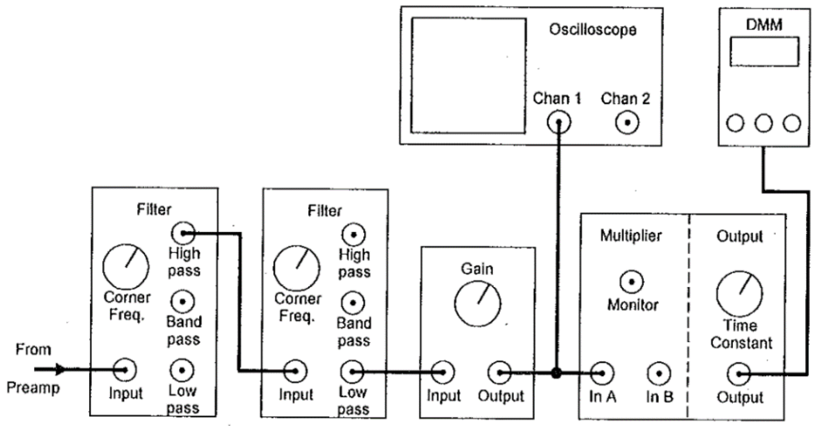

Fig. 4: Cabling diagram for multiplier used as squarer. High—pass filter 0.1 kHz; Low-pass filter 100kHz; Gain 400; multiplier AxA, (all stages a.c. coupled)

As a test further test that the squarer is working, use the XY-display capability on the scope. Connect the squarer's input $V_{\mathrm{in}}(t)$, both to the squarer and to the x-channel of your scope, and connect $V_{\mathrm{out}}(t)$ to the y-channel, and have a look at a real-time xy display. What do you see? 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 the shape on the scope if you raised the gain in the main-amplifier module of the HLE, try it and see? Why does your data lie on a parabola, after all?

Now set the scope back to the normal y vs time setting. Look at $< V_{\mathrm{out}}(t) >$, the time-average of $V_{\mathrm{out}}(t)$. This time average is not zero, why?

Connecting the output of this last stage to a digital multimeter allows you to move easily read $< V_{\mathrm{out}}(t) >$, the needle gauge is just an indicator.

How does the mean value of the meter reading (V_{meter}) relate to the quantity you want to find, namely the mean-square Johnson-noise voltage at the source resistor, $< V_{\mathrm{J}}^2(t) >$?

Why does the meter reading fluctuate, what effect do different averaging time constants have and what is the disadvantage (if any) of just measuring with the longest time constant?

Are there any other contributions to this meter reading? Yes, as you will see in the next section.

1.3 Observing and Correcting for Amplifier Noise

You've now seen how all-analog electronics can take you from a Johnson noise source voltage $V_J(t)$ to a time-averaged d.c. voltage which is a traceable measure of $< V_J^2(t) >$. You will now see 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 $\pm 10\text{V}$. If you were to put a simple sine wave through the amplifiers, you could use the full $\pm 10\text{V}$ excursions. But since you are amplifying noise, you have to ensure that even the rare large fluctuations of the noise stay within the $\pm 10\text{V}$ '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 Fig. 5. For an average noise signal of 3 Volts rms, an excursion beyond $\pm$ 10V is so rare as not to spoil the accuracy of your measurement.

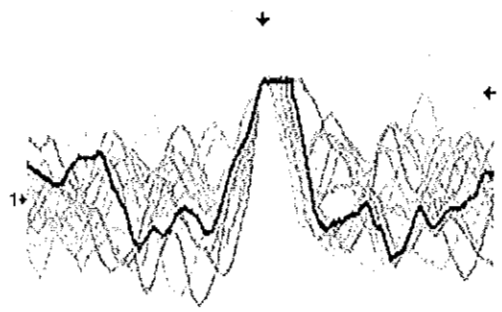

Figure 5: A clipped signal, due to saturating amplifier in HLE

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

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

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

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

so that the time-average at the OUTPUT will be

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

You could use a smaller rms size for the input $V_{\text{A}}(t)$, but you'd be getting an even smaller output from the squarer, and your readings might be affected by the fact the squarer is not a perfect x^2 operation, Think of the "zero offset" you measured in Section 1.2. What is $< V_{\text{sq}}(t) >$ with 0V applied to the squarer input?

From here onwards, whenever you measure a noise voltage, you should check the main-amp output to see that it fits easily into the $\pm$10V range. If it exceeds these limits, reduce the gain. And you should look at the squarer's output on the HLE 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 take a reading of the time-average of the squarer's output, note down the net gain you have used to attain that reading, since this is required to trace the meter reading back to the desired mean-square noise $< V_{\text{J}}^2(t) >$.

b) Now back to Johnson noise. The problem you are now going to address is tracing noise back to its 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’ can be expected to be just as featureless and random as the resistor's Johnson noise, there's apparently no way to separate the two waveforms once they are 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 is a demonstration: let $V_{\mathrm{J}}(t)$ be the instantaneous noise voltage from the source resistor, and let $V_{\mathrm{N}}(t)$ be the instantaneous noise voltage apparently present at the input of the amplifier. That is to say, $V_{\mathrm{N}}(t)$ 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

$$V_{\mathrm{out}}(t) = G\left[V_{\mathrm{J}}(t) + V_{\mathrm{N}}(t)\right],$$

and the mean-square of this output will be

$$< V_{\mathrm{out}}^2(t) > = G^2 < \left[V_{\mathrm{J}}(t) + V_{\mathrm{N}}(t)\right]^2$$
$$= G^2\{< V_{\mathrm{J}}^2(t) > +2 < V_{\mathrm{J}}(t) \cdot V_{\mathrm{N}}(t) > + < V_{\mathrm{N}}^2(t) >\},$$

There is a ‘cross term’ in this expression, the time average of the product $V_{\mathrm{J}}(t) \cdot V_{\mathrm{N}}(t)$, but this time average is zero. The reason is that $V_{\mathrm{J}}(t)$ and $V_{\mathrm{N}}(t)$ can be safely assumed to be uncorrelated, arising as they do from distinct physical mechanisms in two different objects. So when $V_{\mathrm{J}}(t)$ happens to be positive, the amplifier noise $V_{\mathrm{N}}(t)$ 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 is why the absence of correlation enforces a zero for the time- average of the product. Therefore

$$< V_{\mathrm{out}}^2(t) > = G^2\{< V_{\mathrm{J}}^2(t) > + 0 + < V_{\mathrm{N}}^2(t) >\},$$

which says that mean-square voltages from uncorrelated sources are simply additive – like

R_{in} chosen	Gain G_2 (HLE)	$< V_{\mathrm{sq}} >$ read	$< V_{\mathrm{J}}^2 + V_{\mathrm{N}}^2 >$ inferred	$< V_{\mathrm{J}}^2 >$ derived
1Ω	1500	0.6353 V	7.843×10^{-12} V^2	$\sim 0.002 \times 10^{-12}$ V^2
100Ω	1500	0.6516	8.044	0.203
$10\mathrm{k}\Omega$	1000	0.9801V	27.225	19.384
Note these are example values yours will be similar but not the same.				

independent errors “adding in quadrature”. 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 $< V_{\mathrm{J}}^2(t) > = 0$, but in practice, it is good enough to use the $R = 1\Omega$ or 10Ω settings for giving a $< V_{\mathrm{J}}^2(t) >$ which is small enough that the result is a good measure of $< V_{\mathrm{N}}^2(t) >$. Once this amplifier noise value is measured, it can be assumed to be always present, and unchanged, in any use of (the some configuration of) the amplifier. This assumes negligible op-amp current noise, and no noise from external interference, both of which may actually depend on R.

Create a table like that shown above and **plot the data as you take it**. After you have a number of R_{in} and $< V_{\mathrm{J}}^2 + V_{\mathrm{N}}^2 >$ values you should be able to extrapolate to 0Ω to find $< V_{\mathrm{N}}^2 >$ and hence finally derive estimates of the mean-square Johnson noise of the source resistor, corrected for the effects of amplifier noise. You should now correct for amplifier noise in all future measurements in the same way.

Notice that the amplifier-noise corrections are large, even dominant, for small values of source resistance. At what R_{in} are the amplifier and Johnson noise comparable?

Note: when calculating $< V_J^2 + V_N^2 >$ you will need to know both G_1 and G_2 (and don't forget the factor 10V.)

1.4 Johnson noise dependence on resistance

In the previous sections you have seen 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, $< V_J(t) >$, for any source resistor from $R = 10\Omega$ upwards. You should now investigate systematically the dependence of $< V_J(t) >$ upon source resistance R. To do so, you can use the $R = 10\ \Omega$ to $1\text{M}\Omega$ choices built into the pre-amp module. (You can assume these internal source resistors have tolerances of 0.1% to 1 MΩ , and 1% thereafter.)

The three extra positions on the selector switch, A, B, and C, can have other resistors connected. To change these you need to unscrew the four thumb screws on the LLE box and flip over the front panel to expose the back (component) side of the pre-amp's circuit board. You will find the pre-amp power switch (near the internal power-on red LED) inside the low-level electronics), and should turn OFF the pre-amp power before making any changes to the board.

Find the screw-contact terminal strips and check or swap the resistors in positions A, B, and C.

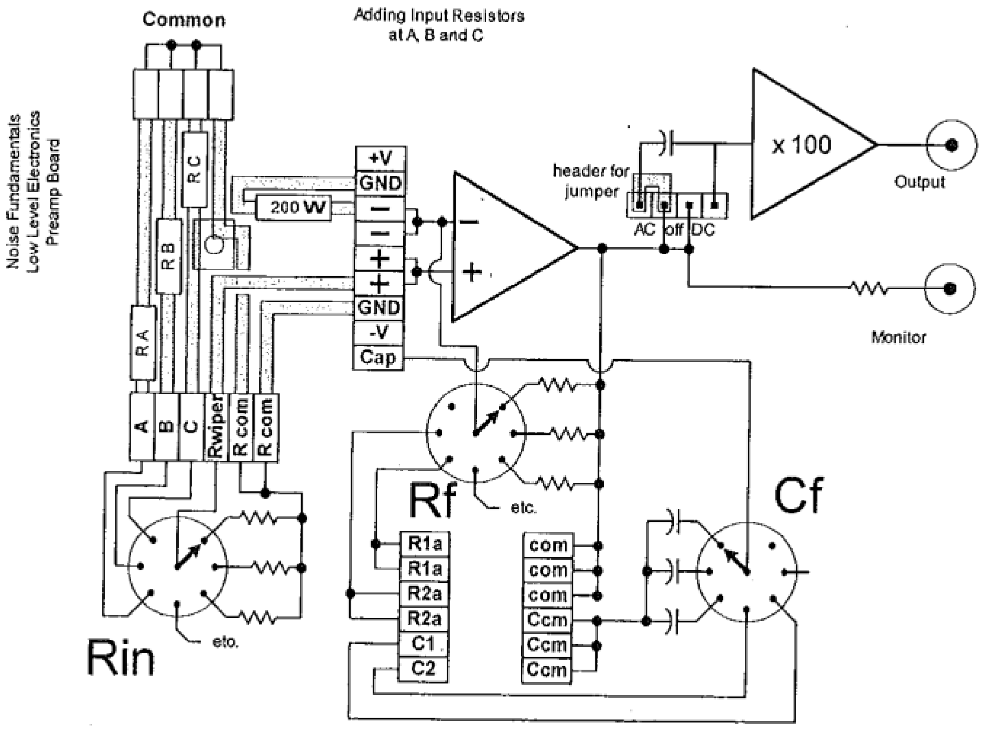

Figure 6: Wiring diagram for adding components at the A, B, C, positions of the pre-amp's input. Note all input resistors have a common ground.

You can try a variety of resistors and take noise data for these resistors, as well as for the built-in source resistors. **As you measure** $< V_J^2(t) >$ values, each corrected for amplifier noise**, you should plot** these as a function of R. Since both axes will vary over many orders of magnitude, a log-log plot may be appropriate.

What are the units on the x and y axes? Nyquist's theory predicts a first-power power-law dependence on resistance R, namely

$$< V_J^2(t) > = 4k_B T \Delta f \cdot R^{\alpha} \qquad \text{where } \alpha = 1.$$

To test whether this prediction is consistent with your data you will also need error estimates on your graph. There will be deviations from this behaviour at the high-R end of the plot, for reasons you will look at next.

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, what does this do to the uncertainties on these points. This would be a particularly good point to demonstrate appropriate use of error bars!

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 i.e. on the range of frequencies to which your system is sensitive. You should now investigate to see if this fits the experimental measurements at a fixed R-value (10kΩ is sensible choice) how the choice of bandwidth matters. The method is to imagine a 'white noise spectrum', i.e. noise power uniformly spread in frequency at its origin, but subsequently modified by the high-pass and low-pass filter sections in the HLE box.

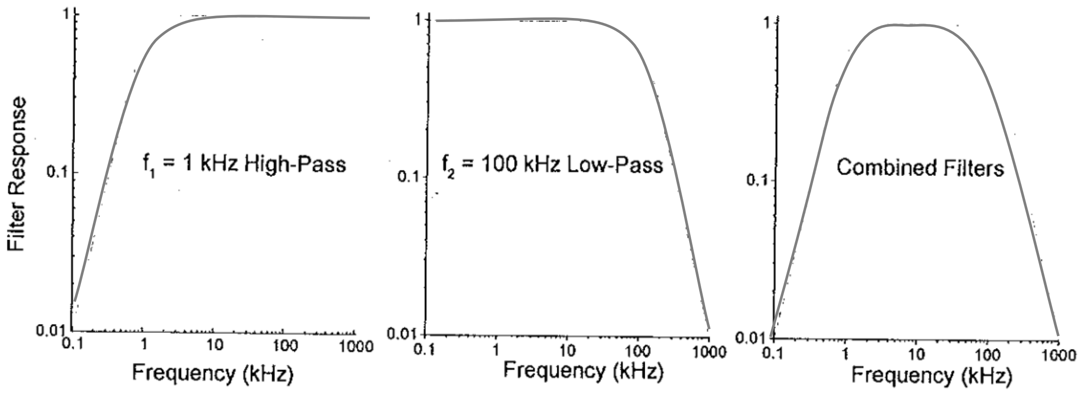

Figure 7: Representation (left) of the transmission of a high-pass filter, of corner frequency f_1; (center) of a low-pass filter, of corner frequency f_2; (right) the combined effect of both filters. Graph axes are logarithmic

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 may at first think that the bandwidth Δf should be given by $|f_2 - f_1|$, which is a decent approximation, but subject to corrections. For now assume $\Delta f = |f_2 - f_1|$.

Measure the mean-square Johnson noise of the resistor, $< V_J^2(t) >$ for a range of (f_1, f_2).

Make three plots of for your data
a function of changing f_1, (const f_2)
a function of changing f_2 (const f_1)
a function of $f_2 - f_1$.

ENBW	$f_2 = 0.33$kHz	1kHz	3.3kHz	10kHz	33kHz	100kHz
f_1 = 10Hz	355	1100	3654	11096	36643	111061
30 Hz	333	1077	3632	11074	36620	111039
100 Hz	258	1000	3554	10996	36543	110961
300 Hz	105	784	3332	10774	36321	110739
1k Hz	9	278	2576	9997	35543	109961
3k Hz	0.4	28	1051	7839	33324	107740

These computed values are all subject to uncertainties of order 4%

Also try plotting it as a function of the equivalent noise bandwidth (ENBW) given in the table above. Which plot is the most nearly linear?

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 typically denoted by S.

1.6 Johnson noise density, and Boltzmann's constant

You have now seen how to measure noise, and have tested its dependence on source resistance and measurement bandwidth. This section introduce noise density, and relates your measured values, via Nyquist's formula, to the Boltzmann constant.

If you have found that the measured mean-square noise $< V_J^2(t) >$ has a linear dependence on the bandwidth Δf, then you have a 'noise density' that is uniform in frequency. Here's an analogy to mass density that should make this clear. Suppose you have a string, of unknown composition, laid out along the 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 it) 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 the mean-square noise $< V_J^2(t) >$ is always proportional to the bandwidth Δf you used to measure it, then you can define the 'noise power density' with a single number

$$< V_J^2(t) > / \, \Delta f,$$

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 $V^2(t)/R$ is 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 $R = 10\ \mathrm{k\Omega}$ 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. How does your data compare to Nyquist's formula written as

$$\text{noise density } S = < \mathrm{V}_J^2(t) > / \Delta f = 4 k_B T R \,.$$

So you should plot all of your data thus far for various R- and Δf-values, to see if you can further establish the linear-in-R claim of the prediction above. Do you see deviations from linear behaviour at large or small R and Δf? Come back to this point after doing Sections 2 and 3. Are the deviations from linear behaviour explainable now?

If you establish a regime of linear dependence on R, your plot will give you a value for a slope, ($4\ k_B T$). What value does it have, (with units)? What uncertainty do you assign to your value?

Now measure the room's temperature T to find a value (and uncertainty) for Boltzmann's constant.

2. Noise Density

2.0 Setting up to see a bandwidth

A particular features of noise measurements (compared to other measurements you may have previously made) is that the signal measured depends on the bandwidth. You have seen earlier that the amount of Johnson-noise (in the 'mean square' sense) depends on the choices made for the difference between f_1 and f_2, the high-pass and low-pass corner frequencies. Here you will ignore the Johnson noise for a while and concentrate on the depiction of bandwidth.

You should use the signal generator to drive the set of filters and plot a graph showing the gain-vs-frequency profile $G(f)$.

A suitable cabling diagram for the system is shown in Figure 8. It requires only the filter sections of the high-level electronics.

A suitable first choice of filter settings would be f_1 = 1 kHz and f_2 = 10 kHz. You can read the peak to peak height of the sine waves from the scope screen or use the built in measurement facility. Make sure you understand what the scope is measuring and what settings are suitable.

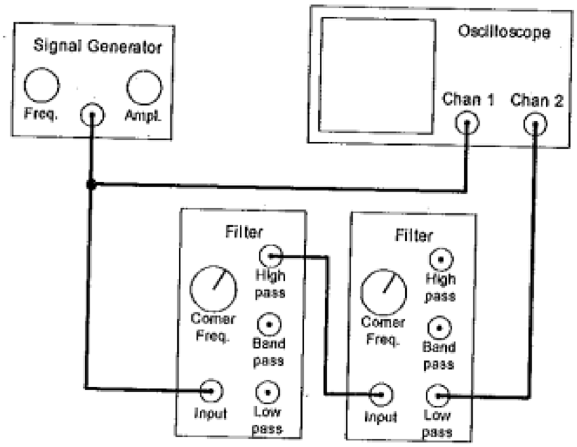

Figure 8: A cabling diagram for testing the action of combined filters on a test signal,

Now for any particular frequency, you can define the gain of the filter assembly as

$$G(f) = [\text{filter output signal}] / [\text{filter input signal}] .$$

For an arbitrary choice of f read the scope and calculate $G(f)$. Is triggering on Ch1 the sensible choice, why?

Notice that for a generic choice of frequency, the output signal will not be in phase with the input. Any "real-time" filtering system will have such phase shifts of output relative to input. In some experiments these phase shifts are important however here they are just a curiosity (and not part of the $G(f)$ definition above). Before moving on, sweep the frequency and observe how the phase shift changes as you go from below f_1 to above f_2. In words, or with a brief hand drawn sketch, note down the behaviour you have seen – no numbers required but LABEL YOUR AXES! (if you sketch something).

It is also possible to measure the gain using a standard digital multimeter (DMM) to perform these rms measurements. Typically the specifications of true-rms a.c. voltmeters extend to 300 kHz, but you would need to check the specifications of your own meter. Briefly compare the scope and DMM V_{rms} readings at low ($\sim f_1$), high ($\sim f_2$) and very high frequency – you only need to note anything down if the expected behaviour is not seen.

Now measure $G(f)$ as a function of f for filter settings of f_1 = 1 kHz and f_2 = 10 kHz. You should cover at least the range of 0.1kHz to 100 kHz.

For (at least) one frequency, change the amplitude of the input signal and check if the output amplitude changes in proportion. This is (one) test that the filter is linear.

In your plot, identify the high-pass corner near f_1 and the low-pass corner near f_2 . Also identify the 'pass band' as the region in which the combined filter assembly gives $G(f) \sim 1$.

You have used sine waves as test signals to be injected into your filter assembly, but is this valid for noise waveforms which are clearly not sine waves? The answer comes from the linearity property, which you tested (for single-frequency sine waves) above. Systems which are linear (or linear within a certain range) display another linearity property, in that:

their response to a sum-of-inputs is equal to

the sum of their responses to the inputs taken individually.

This is the same idea that allows Fourier decomposition. You can think of a noise waveform as being made up of (or, as analysable back into) a whole collection of sine waves, of all frequencies ranging from f << 1 Hz to f >> 1 MHz. The basic part of the linearity property is to understand the operation of the filter as

filter's output (when driven by noise)
= filter's output (when driven by a sum of sine waves)
= sum-of-(filter's output when driven by individual sine waves)

That is to say, the filter's effect on a sum-of-inputs is the same as the sum of the (filter's effect on individual inputs). And the filter's response to individual sine waves is just what's described by the $G(f)$ function you've already measured. Now it is clear that raw noise is composed of all frequencies, but that filtered noise has had its frequency content well below f_1, and frequency content well above f_2, suppressed. Frequency content in the range $f_1 < f < f_2$ is passed along with gain about 1, but frequencies outside the 'pass band' suppressed.

If the 'edges' of the filter's response curve were perfectly sharp-edged corners, you would get a "brick wall" model

gain factor = 1 for $f_1 < f < f_2$, but gain = 0 elsewhere.

and the filter bandwidth Δf would clearly be given by $\Delta f = |f_2 - f_1|$. As you have seen in your $G(f)$ plot, real filters do not have such sharp-edged characteristics. Like many filters the one used in the HLE box are optimized for predictability of performance, rather than sharpness of edge.

It is convenient to approximate a more complicated smooth filter function as a simple rectangular pass band, the "brick wall" model and adjust the width of this model filter function so the integrated area under the gain curve match the real filter, see Fig 9.

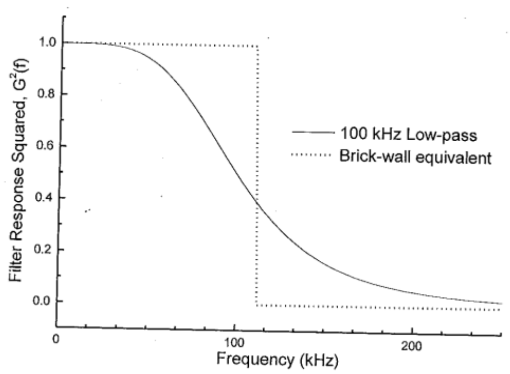

Figure 9: $G^2(f)$ for a typical low pass filter (solid line) and brick-wall filter response of equivalent bandwidth (dashed line).

Looking at the gain curve of the filter from your last plot you can see it has a long 'tail' extending far above the nominal corner frequency. *You may need to add a few frequencies to fill in gaps in your curve at this point.* Calculate/Estimate the area under your gain curve (note log plots do show areas equally) and so the ENBW for the filter given by f_1 = 1 kHz and f_2 = 10 kHz.

Approximately what fraction of the area lies outside this ENBW? What is the fraction contained in the range $f > 4 \times f_2$ and $f > 8 \times f_2$?

What might you conclude about the performance of circuit components well outside the intended frequency measurement window?

3. Johnson Noise comparison at two temperatures

In this last section you will compare $< V_J^2(t) >$ at two different temperatures, room temperature and liquid nitrogen temperatures. It is not practical to put the LLE box in nitrogen but an extension probe is provided with the same resistors as the LLE.

Reconfigure the connections inside the LLE so that the ABC positions for R_{in} of the preamp module connect to the A, B, C terminals of the temperature module. Connections in the variable temperature probe can be checked using the small breakout box. A temperature measurement diode and heater (both not used here) are connected at $D_{1/2}$/$H_{1/2}$. The three resistors are labelled accordingly.

The resistors at the end of the probe are further away from the initial preamplifier and filtering circuits, this means longer cables which can pick up more external noise. The probe is carefully screened with a continuous metal shield but this adds capacitance between the resistor and preamplifier.

Compare some rms noise values at a range of Δf bandwidths (both small and large) to see the effect (if any) of the cable. If the added capacitance due to the probe was around 100pF at what frequency would you expect to notice a difference? Will measurements of all the resistors be affected equally? How would you suggest modifying the experiment to overcome or avoid this capacitance effect?

When you have performed all your room temperature measurements (what is room temperature?) ask the tutor to collect nitrogen and fill the dewar then lower the probe into the glass dewar of nitrogen. It is only required to dip the thin brass spike into the nitrogen, (liquid nitrogen is not conductive but some of the components may not like the sudden thermal contraction if rapidly dunked.) After a relatively short time the bottom of the probe with all be at the same temperature as the liquid nitrogen, what temperature is this?

Please note liquid nitrogen is obviously extremely cold and the probe dipped into it will also become very cold. Be particularly careful of metal parts to which your skin can become instantly stuck resulting in serious cold burns or frostbite. Freezing the water in most cells causes a lot of damage!

Do your experimentally determined cold noise values fit the predicted temperature behaviour for Johnson noise, or more accurately the question should be “does the theory accurately describe the experiment” and are any differences between your values and theory fit significant?

* now return to the question at the end of section 1.6. Do the deviations from linear behaviour you observed in the noise density S make sense, explain your thoughts.