31 A Short Introduction to Noise

v1.9– June 2021

Prerequisite knowledge required: Ohm and Kirchhoff's Laws, Op-Amps, Electrons in Solids

31.1 Introduction

Noise (by which I mean any unwanted signal that appears superimposed on a wanted signal) is always present in electronic systems: all real electronic circuits and components produce noise and all signals in the real world contain some amount of added noise (see figure below). While it can never be entirely removed, it can be quantified and minimised, and the techniques required to do this are vital knowledge for any engineer wanting to get the best performance from circuits.

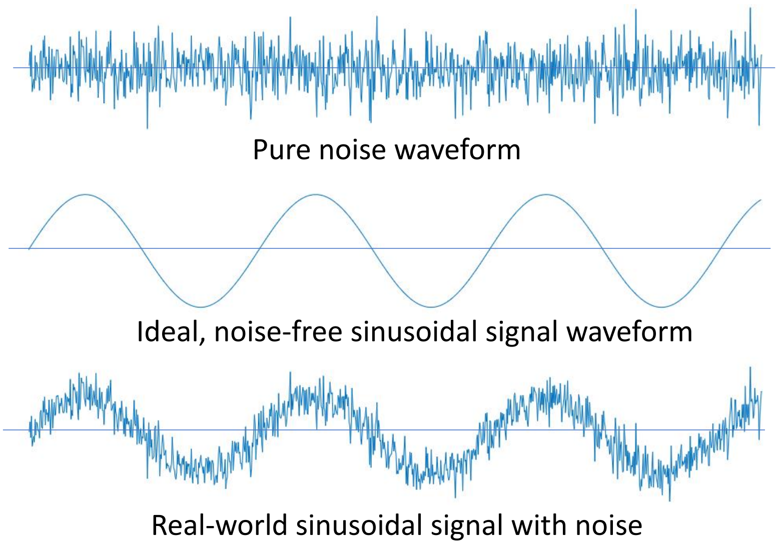

Figure 31.1 Typical waveforms of noise, ideal sinusoidal signal, and real-world sinusoidal signal

Noise can be divided into three main categories depending on where it comes from:

1. *Intrinsic noise*: noise generated within the elements of the circuit due to the random motion of the charge carriers.
2. *Extrinsic noise*: noise generated externally to the circuit or system under consideration and coupled into the circuit usually through electric or magnetic fields.
3. *Quantisation noise*: noise introduced onto a signal due to the operation of an analogue-to-digital converter.

We'll be considering all three types of noise in this chapter: finding out how the noise arises, how to minimise the noise levels in a circuit, how to treat noise as a mathematical signal, and how to measure the noise levels in a circuit.

31.2 Some preliminary statistics: Gaussian and white noise

(Note that this section is not essential for understanding of the rest of this chapter. If you haven't studied statistics before and find this hard-going, please feel free to jump ahead to the next section.)

The probability density function *p(x)* of a statistical variable evaluated at *x* is a measure of how likely the variable is to have a value close to *x*. To define it more exactly: the probability of any sample of the noise lying between the values of *x* and *x + dx* where *dx* is a small range is *p(x)dx*, and for a wider range, the probability of *x* lying between *a* and *b* is:

$$prob(a \le x < b) = \int_a^b p(x)\,dx \tag{31.1}$$

It follows that the integral of any probability density function from minus infinity to plus infinity is one, since the probability that any sample of the variable has a value between minus and plus infinity is one: it must always be somewhere in that range.

The two most important intrinsic noise sources are often described as being *white*, *Gaussian* noise sources. The "Gaussian" part of this means that the probability density function follows a Gaussian (sometimes called a "normal") distribution, with a standard deviation of σ:

$$p(x) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{x^2}{2\sigma^2}\right) \tag{31.2}$$

The standard deviation is a measure of the power in the noise: the mean squared value of a noise waveform with a normal distribution given by equation (31.2) is the square of the standard deviation.

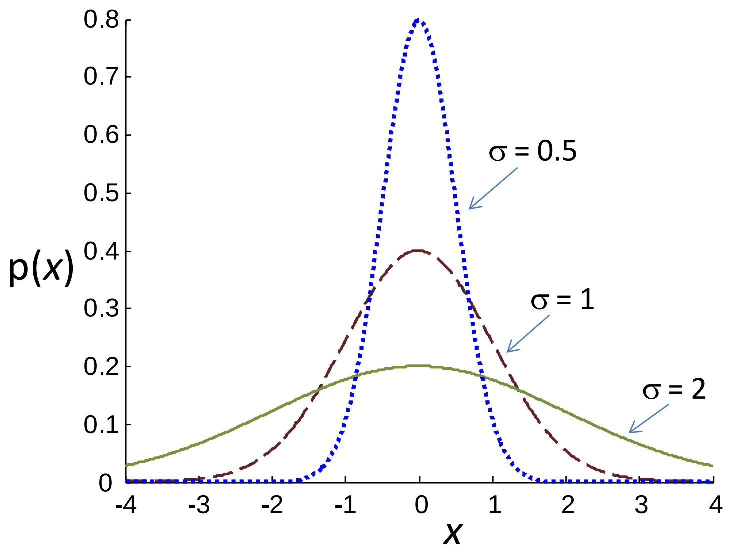

Figure 31.2 Zero-mean Gaussian amplitude distributions with different standard deviations

Gaussian noise is very common because of the Central Limit Theorem[1]. For example, intrinsic noise is caused by very large number of individual noise contributions (electrons) moving randomly

[1] If you've not come across the Central Limit Theorem, it says: "the probability distribution of the sum of a very large number of independent random variables has a Gaussian distribution". It doesn't really matter what the probability distribution of the individual random variables is, as long as there are enough of them the probability distribution of the sum of them approaches a Gaussian distribution. In the case of electrons moving around randomly in a resistor, the total noise voltage is caused by the sum of the movements of all of the electrons, which satisfies this criterion.

around, and this tends to result in a Gaussian probability distribution, independent of what the probability density function might be for any individual electron.

The standard deviation of the noise distribution is related to the mean power in the noise by:

$$\text{Mean Noise Power} = \frac{\sigma^2}{Z} \tag{31.3}$$

where Z is the system impedance (if no system impedance is specified then Z is usually assumed to be equal to one, and the power becomes a normalised power). This means that the standard deviation of the noise probability density function is equal to the rms amplitude of the noise waveform (this is true for any noise distribution with a zero mean). This is one reason why the rms amplitude of a signal is so useful: it is simple to calculate the power in the signal from the rms amplitude: you don't need to know anything else about the shape of the waveform.

Noise can also be quantified in terms of its power spectral density (PSD), which is a function of frequency, and is defined in a similar way to the probability density function: $PSD(f)df$ is the power in a small range of frequencies df around f, and the power in a noise signal in range of frequencies between f_1 and f_2 is given by:

$$\text{power in frequency range } f_1 \leq f < f_2 = \int_{f_1}^{f_2} PSD(f)\,df \tag{31.4}$$

The total noise power in a range of frequencies from zero to B Hz is therefore given by the integral of the noise power spectral density over the range of frequencies of interest:

$$\text{Mean Noise Power} = \int_0^B \text{PSD}(f)\,df \tag{31.5}$$

(Essentially, this is the noise power that would emerge from an ideal bandpass filter which lets through any frequencies between A and B, and no others.)

As well as having a Gaussian (normal) amplitude distribution, the main intrinsic noise sources can usually be assumed to be *white*, which means that the noise power spectral density at all frequencies is the same[2]. When the noise is white and hence the power spectral density of the noise is constant, equation (31.5) reduces to:

$$\text{Mean Noise Power} = B \times PSD \tag{31.6}$$

Whilst most noise when it is first generated is approximately white, this is not true after the noise has passed through a frequency-dependent stage such as a filter or an amplifier with a limited

[2] This isn't exactly true, since it leads to the conclusion that all noise has an infinite power (there are an infinite number of frequencies, so if the noise spectral density really was constant, then the integral of the noise spectral density across all frequencies would also be infinite, and that implies an infinite noise power, which is clearly silly). In practice the noise is white for all frequencies of interest to most electronic engineers, so it's a safe approximation here to treat the noise as white.

bandwidth. However, by using the concept of an "equivalent noise bandwidth", we can always express the output noise from such a system in the form:

$$\text{Mean Output Noise Power} = Gain \times B_N \times PSD \tag{31.7}$$

where *Gain* is the gain of the system, B_N is the equivalent noise bandwidth of the system, and *PSD* is the constant power spectral density of spectrally-white input noise. (The equivalent noise bandwidth is the bandwidth of a brick-wall[3] filter, which if placed after a perfect amplifier with the same gain would let through the same total noise power. For a first-order (single-pole) filter the equivalent noise bandwidth is π/2 times the 3-dB bandwidth[4].)

Note that all of these equations deal with noise power, not noise voltage. To get the rms noise voltage you have to take the square root of the mean noise power, which means multiplying the square-root of the power spectral density by the square root of the bandwidth:

$$\text{rms Noise Voltage} = \sqrt{B} \times \sqrt{PSD} \tag{31.8}$$

Sometimes the noise is specified in terms of “volts per root Hz”. This is actually a measure of the square root of the power spectral density: multiply the noise in V/√Hz by the square root of the bandwidth (which has units of √Hz) and you get a quantity with units of volts: this is the rms noise voltage.

31.3 Intrinsic noise

Intrinsic noise in electronic circuits (sometimes also called circuit noise) can itself be divided into several categories, the most important of which[5] are:

1. *Thermal noise*: noise generated by thermally induced random motion of charge carriers in all conductive elements. It is present even in the absence of current flow. Thermal noise is dependent on temperature and component value: noise increases at higher temperatures, and larger resistances produce higher noise voltages.
2. *Shot noise*: noise resulting from the random flow of charge across a potential barrier, e.g. across a semiconductor junction. The power of the shot noise is proportional to the current flow. Larger currents provide more shot noise, but the ratio of shot noise to current decreases as the current increases.
3. *Flicker noise*: noise which exists mostly at low frequencies; it tends to decrease in power inversely proportional to the frequency, hence its other common name: “1/f noise”.

31.3.1 Thermal noise

Every resistor in every electronic circuit (or even one just lying on the lab bench not in a circuit) has a continuously changing voltage across it due to the thermal movement of the electrons within it. They all move (partially) independently, so there is a chance that a large number of them might

---

[3] A brick-wall filter has a gain of one for all frequencies in the passband, and a gain of zero for all other frequencies. They don’t exist, but they are a useful mathematical model of an “ideal” filter.

[4] It’s a good exercise to derive this result from the frequency response of a first-order filter.

[5] There are a couple of other intrinsic noise sources, including the entertainingly called “popcorn” noise (or burst noise) which sounds like popcorn popping if amplified.

decide to move up one end at the same time, and that causes a higher negative charge at one end than at the other, and the result is a non-zero electric field, and hence a voltage across the resistor.

It's usually modelled as white (since it has a constant spectral density up to frequencies far in excess of those we're interested in here) and Gaussian. For circuit analysis, it can be modelled as a small voltage in series with the resistor, the voltage having a mean-square magnitude of:

$$\overline{v^2} = 4kTBR \tag{31.9}$$

where k is the Boltzmann constant (about 1.38e-23), T is the absolute temperature (usually assumed to be 290 K in the lab[6]), B is the bandwidth in Hz, and R is the value of the resistor. Note that the larger the resistor, the greater the noise voltage across its terminals[7].

This means that any real-life (noisy) resistor can be modelled by the series combination of an ideal noiseless resistor of the same value, and a noise voltage source; the voltage source has a zero-mean Gaussian (normal) probability distribution, and an rms value (equal to the standard deviation of the noise distribution) given by:

$$e_t = \sqrt{4kTBR} \text{ volts} \tag{31.10}$$

The symbol e_t is used since this is thermal noise.

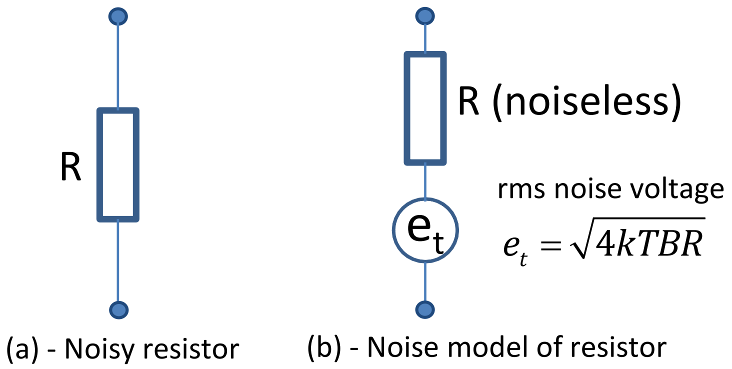

Figure 31.3 Resistor noise model

Note that this suggests the use of smaller resistors to achieve the lowest noise voltages. It also predicts that the colder the circuit, the less of this form of noise there is. (This is the reason why extremely sensitive radio receivers are cooled with liquid nitrogen.)

---

[6] Why 290 K? That's about 17 degrees Celsius, which is a bit colder than most labs. However using this value means that kT is almost exactly 4.0 x 10^{-21}, which is easy to remember and use, and it's close enough for most purposes.

[7] This results in an interesting puzzle: the formula predicts that the noise voltage across an infinite resistor would be infinite. But there is effectively an infinite resistance (nothing) between any point in a circuit and ground. Therefore, there is an infinite noise voltage at all points in all circuits. What's wrong with this argument? (Hint: think about potential dividers and the output impedance of the noise source.)

31.3.2 Shot noise

Currents have noise associated with them too. Electrons don't march along in orderly ranks, with exactly the same number passing across any boundary every second. Even when the current is nominally constant, in some seconds a few more than average pass by a given point in a circuit, in some other seconds a few less than average pass by. This effect results in a noise source associated with every current passing across a boundary (for example across a diode). It can be modelled as a small noise current source in parallel with every such current, with a value of:

$$\overline{i^2} = 2IeB \qquad (31.11)$$

where *I* is the mean current, *e* the charge on the current carrier (usually an electron), and *B* the bandwidth in hertz again. Although this noise does increase with increasing current, it's interesting to note that the square of the ratio of the noise current to the signal current (which is often what we're more interested in since it contributes to the signal to noise ratio) is:

$$\overline{\left(\frac{i}{I}\right)^2} = \frac{2eB}{I} \qquad (31.12)$$

so the noise relative to the signal actually decreases with larger currents. Again, smaller resistors (which would produce larger currents) help to minimise the noise.

Unlike thermal noise, shot noise is independent of temperature, so cooling a circuit doesn't help reduce this sort of noise.

31.3.2.1 Why don't you get shot noise in resistors or wires?

There's an often-overlooked point here that's quite interesting: why does shot noise only appear in devices like diodes, where the electrons are jumping from the conduction band down to the valence band? Why doesn't it happen in all currents, for example those in wires? After all, electrons are moving along those in a rather random, chaotic way too.

The answer is that there's another effect in wires which effectively cancels out the shot noise. Imagine, just due to random chance, that a larger than average number of electrons happen to have moved along a small segment of wire in one small instant. The result is a net negative charge at the end of the segment, and this tends to repel other electrons, reducing the number of electrons that move along the same segment in the next small instant.

This smaller number of electrons in the next instant means that over both instants taken together, the larger number of electrons moving in the first instant is compensated by the smaller number in the second instant, making the total the average number expected for the current flowing.

This compensation happens so fast that over the timescales of interest in electronics, this effect tends to practically remove all the shot noise except in situations where the number of electrons moving doesn't significantly change the potential difference that the electrons have to move across, and the most common example of that is a p-n junction where the potential difference is set by the bandgap (the energy difference between the valence and conduction bands).

31.3.3 Flicker noise

In addition to thermal and shot noise, most semiconductor devices exhibit a form of low-frequency noise called "flicker noise" (or sometimes "1/f noise" since the noise power tends to decrease as the frequency increases).

This is why a device like the TL071 op-amp specifies noise as: "V_n = 18 nV/√Hz at 1 kHz". The noise from a device like this is not white: although the spectral density is fairly constant at higher frequencies, at lower frequencies there is a lot more noise (see Figure 31.4).

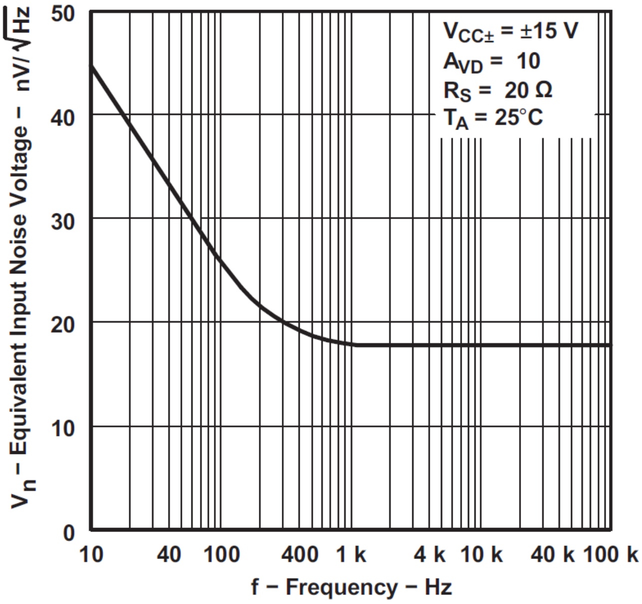

Figure 31.4 Equivalent input noise spectral density for the TL071 op-amp (taken from the TL071 data sheet)

31.4 Extrinsic noise

Extrinsic noise can also be divided into several categories, depending on how it couples onto the signal:

1. *Power supply noise*: noise that has arrived on a signal due to a noisy (non-constant) power supply. Amplifiers will specify a power-supply rejection ratio (the ratio of the noise on the power supply to the noise on the output signal, usually in dB) which can be used to rate the susceptibility of the amplifiers to this sort of noise.
2. *Hum*: noise resulting from mains pickup, usually related to different ground voltages at different points in the circuit.
3. *Crosstalk and pickup*: noise coupling into the circuit from an adjacent external circuit, either by electrostatic or magnetic induction.
4. *Radiated noise*: noise from a distant source that has arrived due to parts of the circuit behaving as a radio antenna.

It's difficult to predict the levels of extrinsic noise, since it depends on the behaviour of circuits outside the circuit you're designing. What you can do is try and make your circuit insensitive to this noise, by screening (placing sensitive inputs inside a metal box to stop radio waves getting there), using balanced inputs (a good approach for mains hum), and providing sufficient decoupling and filtering on the power supply.

All these forms of extrinsic noise share one property: they are not white. They have characteristic spectra, with peaks at certain frequencies. The spectrum of the noise can be is a useful indication of where the noise is coming from: for example anything which arrives from a digital signal often has a series of peaks at multiples of the digital clock frequency; noise from a USB power supply often appears at multiples of 1 kHz, as there is a lot of digital noise that couples across from the data lines in the USB cable to the power supply line, and the USB connection standard uses a communications protocol based on a 1 kHz frame rate; and noise that originated from the mains power supply usually appears at 50 Hz (at least in the UK).

31.5 Quantisation noise

The third type of noise is produced by analogue-to-digital converters. Since all real analogue-to-digital convertors operate with a finite number of bits, none of them can produce an exact representation of the value of an incoming signal at every point in time. There will always be a slight error, and this error, which changes every time the analogue input is sampled, can be treated as a noise source:

$$\text{ADC output (digitised)} = \text{ADC input (analogue)} + \text{noise} \qquad (31.13)$$

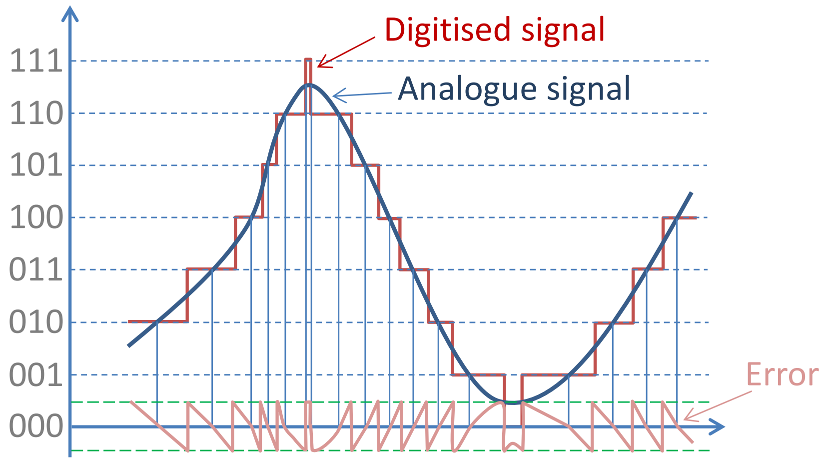

Figure 31.5 Illustration of the source of quantisation noise

The power in this noise can be determined reasonably easily for the case where the possible output (quantised) levels are closely spaced together and the analogue signal never goes outside the range of the analogue-to-digital converter (which is the usual case). In this case, the maximum possible noise voltage is one-half of the difference between two quantisation levels, and the minimum possible is minus one-half of the difference.

For the case of a typical 16-bit ADC and Audacity, the signal is always represented as a number between minus one and one, which means that the distance between two adjacent quantisation levels is:

$$\Delta V = \frac{1-(-1)}{2^{16}} = 30.5\,\mu \qquad (31.14)$$

So the amplitude of the noise always lies between +15.26 µ and -15.26 µ relative to the maximum amplitude, since these are the maximum and minimum possible differences between the input voltage and the quantised voltage.

(It's important to note that this noise, although often assumed to be white, is not Gaussian. The noise voltage cannot take any value; it is always within the range from +15.26 µ to -15.26 µ.)

If the analogue signal source supplying the input to the ADC already has less noise than this, there is little point in improving the noise performance of the analogue circuitry. To get less noise in the digitally recorded signal, it would be necessary to increase the number of bits in the ADC (which means greater storage requirements as well as a more expensive ADC).

In terms of the power in this noise, we can make the reasonable assumption that the noise is equally likely to have any voltage in this range. Then the mean-square value of the noise voltage (which is equal to the average noise power) is:

$$\begin{aligned}\text{mean noise squared} &= \frac{1}{30.52\mu}\int_{-15.26\mu}^{15.26\mu} v^2\,dv = \frac{1}{30.52\mu}\left[\frac{v^3}{3}\right]_{-15.26\mu}^{15.26\mu} \\ &= 77.62\times10^{-12}\end{aligned} \tag{31.15}$$

which corresponds to an rms value of 8.8µ of noise. On top of a signal which can be +/- 1, this might seem a very small amount of noise, however it is often still audible (the human ear is an amazingly sensitive receiver).

It's equivalent to a noise level of:

$$20\log_{10}\left(8.8\times10^{-6}\right) = -101.1\ \text{dBFS} \tag{31.16}$$

where dBFS is decibels relative to the full-scale reading. This is the best that it's possible to do with a 16-bit ADC.

More generally, consider an ADC with N-bits in its output, and an input range from -A to +A volts. The difference between two output levels is now:

$$\Delta V = \frac{A-(-A)}{2^N} = \frac{2A}{2^N} \tag{31.17}$$

and the mean noise squared is:

$$\begin{aligned}\text{mean noise squared} &= \frac{2^N}{2A}\int_{-A/2^N}^{A/2^N} v^2\,dv = \frac{2^N}{2A}\left[\frac{v^3}{3}\right]_{-A/2^N}^{A/2^N} \\ &= \frac{2^N}{2A}\left(\frac{2A^3}{3\times2^{3N}}\right) = \frac{A^2}{3\times2^{2N}}\end{aligned} \tag{31.18}$$

31.5.1 Measuring noise in bits

Since it is possible to determine a noise voltage associated with an ADC with a certain number of bits, it is equally possible to measure a noise voltage in terms of a number of bits. This is sometimes done, and a signal in a system might be specified as having, for example, four bits of noise. This just means that the rms noise voltage is four times the difference in voltage between two adjunct output levels of the ADC.

For a 16-bit ADC with an input range from -1 V to +1V, four bits of noise would correspond to a noise voltage of:

$$V_{rms} = 4 \times \frac{(1 - -1)}{2^{16}} = 4 \times \frac{2}{2^{16}} = 122\,\mu V \tag{31.19}$$

31.6 Testing the noise colour of noise

If you plot the spectrum of noise in a poorly-designed USB audio interface (to find out which frequencies have most noise present), you might get results like those shown in Figure 31.6 (taken using Audacity):

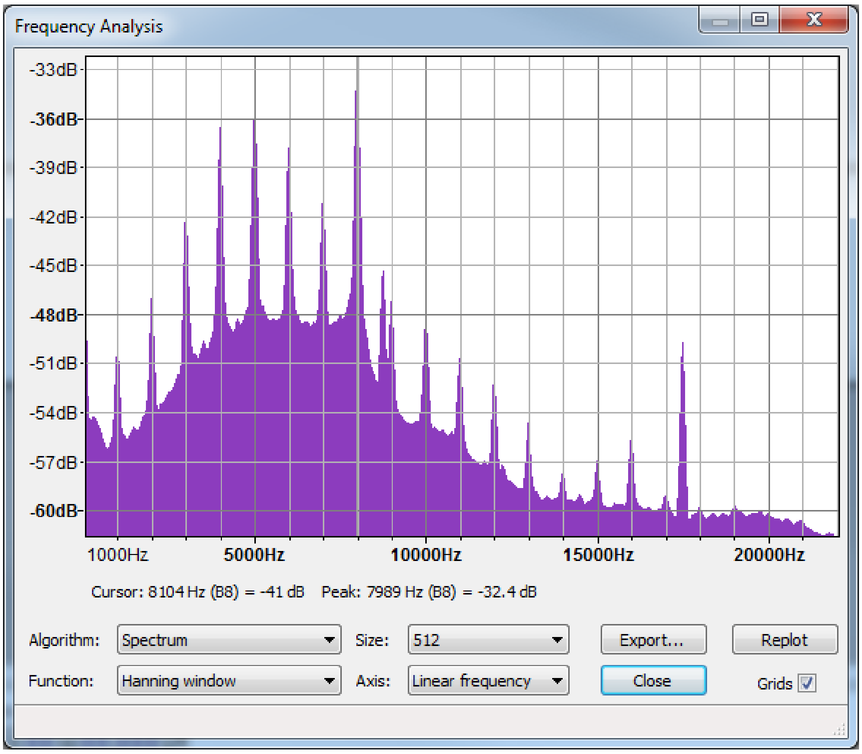

Figure 31.6 Spectrum of Sample Noise Measured Using a VAM v3.1

The units here are dBFS (dB reference to the Full-Scale voltage, although they are labelled as just in "dB") and the points on the graph represent the total amount of power in one "frequency bin".

A "frequency bin" is a range of frequencies of width = sample_rate / 2.0 / Size, centred on the given frequency. For example, the spectrum here has a peak value of -32.4 dB at 7989 Hz, the sample rate is 44.1 kHz (the standard sampling rate used for compact disks) and the Size parameter is set to 512. This means that if a filter[8] was used which allowed through a range of frequencies of width:

[8] I'm skipping over an important detail here: the shape of the filter is important, as well as its bandwidth. You can try out the effects of other filters yourself in Audacity.

$$\frac{sample\_rate}{2.0 \times Size} = \frac{44100}{2.0 \times 512} = 43.07 \text{ Hz} \qquad (31.20)$$

centred around 7989 Hz[9], the rms amplitude of the signal output from the filter would have an rms amplitude 32.4 dB below the full-scale (which as before is 20 x $\log_{10}(1.0) = 0$), so this means that if just this range of frequencies were considered, the total rms amplitude would be:

$$rms = 10^{\left(\frac{-32.4}{20}\right)} = 0.024 \qquad (31.21)$$

relative to the full-scale.

Looking at the spectrum in Figure 31.6 it is obvious that the noise is not white: there are clearly some peaks visible at multiples of 1 kHz. However in other cases, where it might not be so obvious whether the noise is white (i.e. has equal power at all frequencies) or consists of a large number of individual frequency components, one way to investigate is to change the "Size" parameter of the frequency analysis, which changes the size of the frequency bins plotted on the spectrum plot.

For example, increasing the Size parameter by a factor of eight (from 512 to 4096) reduces the size of the frequency bins by a factor of eight. If the noise were white, this would be expected to reduce the total power getting through the filter by a factor of eight as well (and hence the total rms amplitude by a factor of √8 since power is the square of the rms figure).

However if the power around these frequencies is primarily due to a single discrete frequency component, reducing the size of the frequency bin won't make much difference: most of the power will still arrive in the same frequency range.

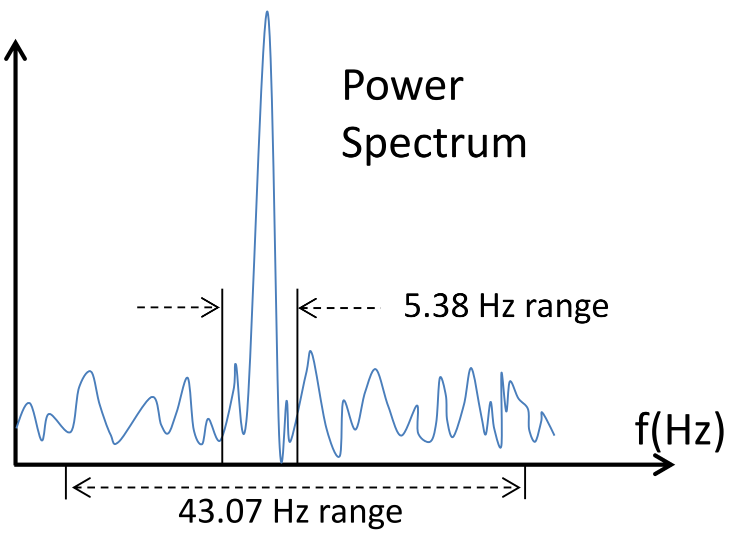

Figure 31.7 Non-White Noise Spectrum with Frequency Bins

[9] It's worth noting that the actual frequency of this noise is probably not exactly 7989 Hz; 7989 Hz is just the middle of the range of frequencies where the noise energy occurs. If the noise is actually coming from a single-frequency source it could be anywhere from 7974.6 Hz to 8101.5 Hz. Given this range, it's a good bet that the noise is coming from a nominally 8 kHz source.

Figure 31.8 below shows the result of analysing the same signal as in Figure 31.6, only using frequency bins one-eighth of the width. As noted above, if the noise was white, then this would be expected to reduce the noise amplitude getting through the filter by a factor of √8 = 2.83 = 9 dB. However in this case the peak value at just under 8 kHz has reduced by only 0.8 dB.

This strongly suggests that almost all of the noise around this frequency is due to a single frequency component, since reducing the size of the frequency range considered when calculating each point does not affect the answer. This in turn points towards an extrinsic source for this noise, rather than intrinsic noise due to thermal, shot or 1/f noise.

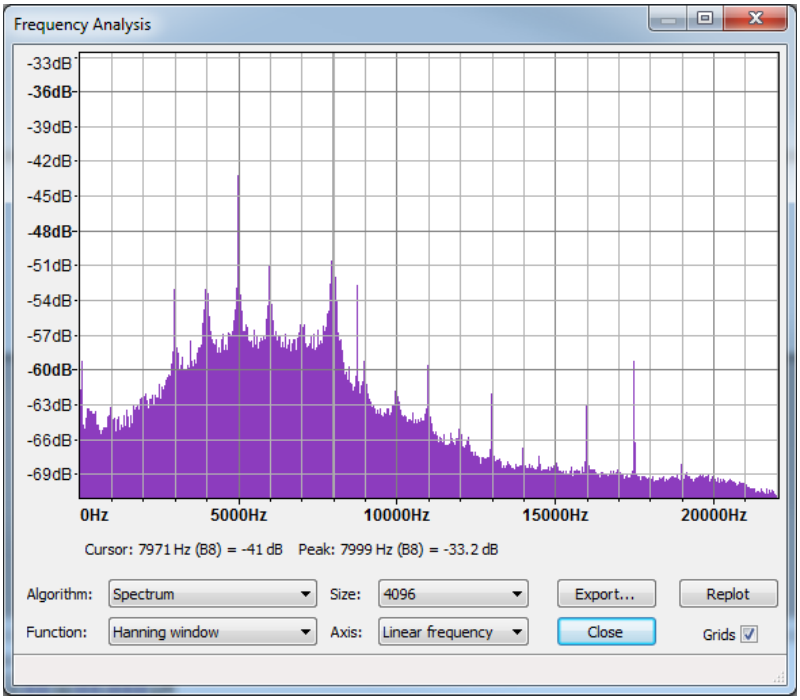

Figure 31.8 Spectrum of Sample Noise using Size = 4096

31.7 Adding noise sources

All of these noise sources have one thing in common: they can almost always be regarded as uncorrelated or independent sources of noise. In other words you can't predict anything about one noise source from knowledge of the current value of any other noise source.

This has one very important consequence for combining two noise sources: consider the mean power in the sum of two independent noise sources $n_i(t)$ and $n_j(t)$:

$$\begin{aligned}\overline{\left(n_i(t)+n_j(t)\right)^2} &= \overline{n_i^{\,2}(t)+2n_i(t)n_j(t)+n_j^{\,2}(t)} \\ &= \overline{n_i^{\,2}(t)}+2\,\overline{n_i(t)n_j(t)}+\overline{n_j^{\,2}(t)}\end{aligned} \tag{31.22}$$

The second term in this expansion is the product of the two noise sources. However, since they are independent, at any time the probability of them having the same sign (and therefore the product being positive) is the same as the probability that they have different signs (and therefore the product is negative). Therefore the mean value of this product is zero: the value is positive the same amount of time that it is negative.

This leads to the equation:

$$\overline{\left(n_i(t)+n_j(t)\right)^2} = \overline{n_i^{\,2}(t)} + \overline{n_j^{\,2}(t)} \tag{31.23}$$

In other words, the mean power in the sum of two independent noise sources is equal to the sum of the mean powers in the individual noise sources.

This is an important point to remember when adding ***independent*** noise sources: ***always add the powers, not the amplitudes***.

31.8 Noise in op-amp circuits

To determine the intrinsic noise introduced by an operational amplifier, a noise model of the operational amplifier must be used, as well as a noise model of the resistance at its input. Since the noise we're most interested in has a uniform spectral density, the noise voltage is often expressed in terms of this constant spectral density: as the noise voltage per $\sqrt{\text{Hz}}$ (pronounced 'root Hertz').

Figure 31.9 represents a general model of an operational amplifier, showing the effect of the noise characteristics usually specified in the datasheets. (In a real op-amp there are many more than three individual sources of noise, as each resistor will have a thermal noise associated with it, and each transistor and diode a shot noise associated with it, but to make noise calculations easier, the contributions of all of these noise sources are combined into these three equivalent input noise sources.)

This is the usual way of specifying the noise added by any amplifier: the total noise due to all of the individual noise sources within the amplifier is transferred to a few equivalent noise sources at the input of the amplifier. In the case of operational amplifiers (which have two inputs) three noise sources are required at the input to specify the noise (however the two noise currents have the same noise power).

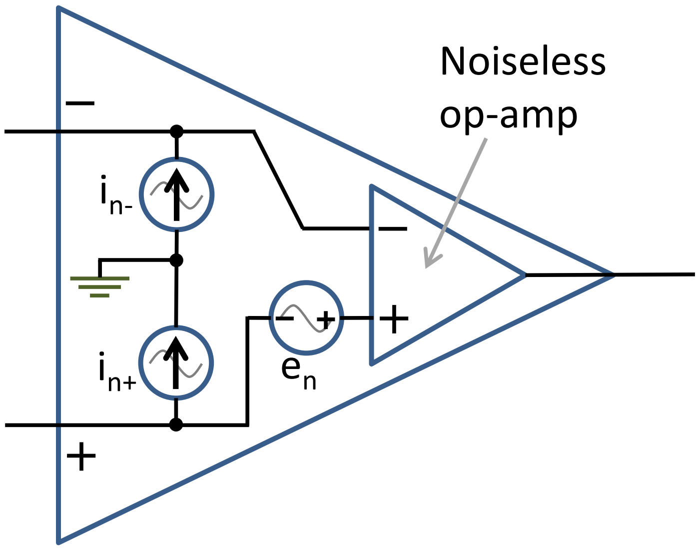

Figure 31.9 General noise model of an operational amplifier

Here e_n is the rms input noise voltage of the op-amp, and i_n is the rms input noise current (note there are two of these, one associated with each input). All noise sources have a zero-mean Gaussian (normal) distribution, and all are independent (so the noise powers from each noise source must be added, not the noise voltages).

The values of e_n and i_n are dependent on the internal design of the op-amp and will be specified in the datasheet. For example, Figure 31.10 shows an excerpt from the datasheet of a TL071 op-amp showing the noise parameters:

V_n	Equivalent input noise voltage	R_S = 20 Ω	f = 1 kHz	18	18	nV/√Hz
			f = 10 Hz to 10 kHz	4	4	μV
I_n	Equivalent input noise current	R_S = 20 Ω,	f = 1 kHz	0.01	0.01	pA/√Hz

Figure 31.10 Excerpt from TL071 datasheet showing the noise parameters

The equivalent input noise voltage e_n is expressed in two ways: first in $\text{V}/\sqrt{\text{Hz}}$ at 1 kHz, this implies that above this frequency the noise voltage source is white, and to convert from the value in $\text{V}/\sqrt{\text{Hz}}$ to the value in volts you can just multiply by the square root of the noise bandwidth at that point in the circuit). Secondly as a fixed 4 μV from 10 Hz to 10 kHz, which suggests that the voltage noise source is not white over that frequency range, so a total noise voltage is given instead. The noise current is only expressed in $\text{A}/\sqrt{\text{Hz}}$, which implies that it can be treated as white at all frequencies.

(Note that both the input noise voltage and input noise current are specified using a source impedance of 20 ohms. This small value of resistance is chosen so that the noise voltage of the input resistor itself is negligible in the measurement[10].)

(It is a common mistake to forget which units a value of noise voltage or current is in and get answers that are wrong by a factor of the square root of the bandwidth.

Another quick reminder: combining the powers, rather than the voltages, is another common mistake when doing noise calculations. The key point to remember when using noise models is that when combining unrelated noise voltages you must add the mean square values[11], as the noise voltages are uncorrelated. For example:

$$e_{Total} = \sqrt{e_{One}^{\ 2} + e_{Two}^{\ 2}} \quad (31.24)$$

Also, ensure that when adding noise voltages they are either both expressed in volts, or both in $\text{V}/\sqrt{\text{Hz}}$.)

31.8.1 Analysing a simple op-amp circuit for noise

Determining the expected noise power at the output of an op-amp circuit is somewhat tedious, since there are a large number of sources that have to be considered. However since all the noise

[10] You can't use a source resistance of zero, since that would imply an infinite gain, and the op-amp would saturate.

[11] The mean square values are just the squares of the rms (root mean square) values.

sources are independent, a slight variation on the principle of superposition can be used: determine the expected noise output voltage due to each noise source in turn, and then combine the noise powers of each contribution to determine the total noise power.

For example, if an op-amp was configured as an inverting amplifier, as shown in Figure 31.11,

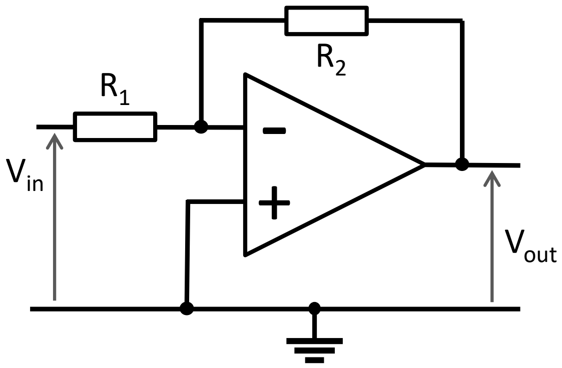

Figure 31.11 Inverting amplifier with non-inverting input grounded

the noise model of this circuit, showing all five noise sources, would be:

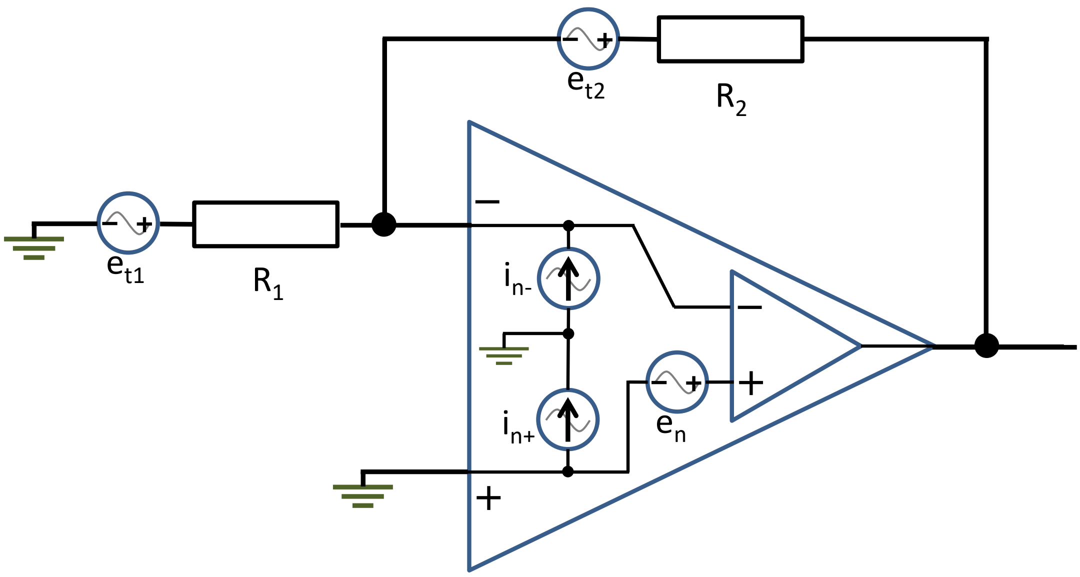

Figure 31.12 Noise model of an inverting op-amp (with non-inverting input grounded)

(The op-amp has three independent noise sources, and the two external resistors both have thermal noise sources associated with them.)

To determine the total noise, we have to work out the effect of each of these noise sources in turn, setting all the other noise sources to zero. First, for the thermal noise associated with the input resistor R_1, the circuit looks like this:

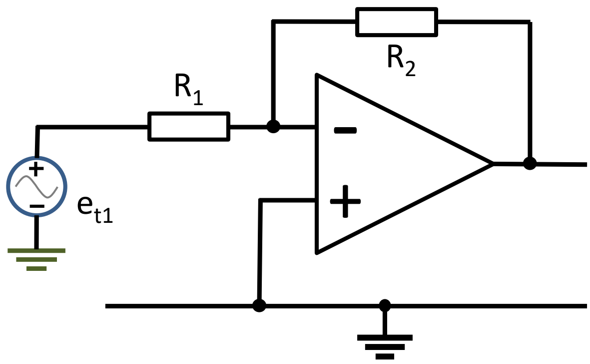

Figure 31.13 Showing the contribution of the input resistor to the noise of an inverting amplifier

and the noise spectral density at the output due to this noise source is clearly:

$$v_{R_1} = e_{t1}\frac{R_2}{R_1} = \sqrt{4kTR_1}\,\frac{R_2}{R_1} = R_2\sqrt{\frac{4kT}{R_1}} \quad \mathrm{V}/\sqrt{\mathrm{Hz}} \qquad (31.25)$$

since this is just an inverting amplifier configuration. (Note that this is a positive result, rather than the negative result usually produced from an inverting op-amp. It's positive for noise since I am working with rms values here, not actual DC levels, so the fact that the output is inverted makes no difference to the rms amplitude; all rms amplitudes are positive.)

Next, for the thermal noise associated with the feedback resistor R_2, the circuit looks like:

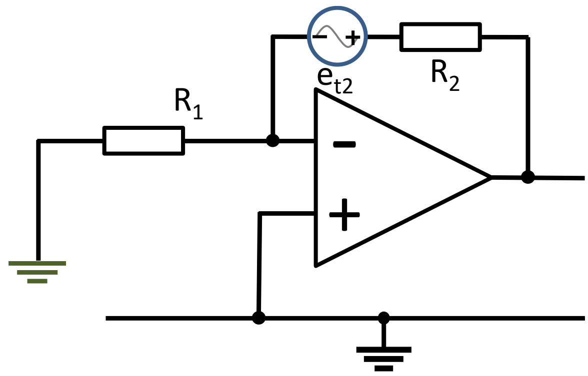

Figure 31.14 Showing the contribution of the feedback resistor to the noise of an inverting amplifier

Here, note that an ideal non-saturating op-amp (with infinite gain and no offset voltage) keeps its two inputs at the same voltage, which in this case means at 0 V (since the non-inverting input is tied directly to ground). There is no current flowing through R_1 (since it is grounded on both sides), and therefore no current through R_2 (as ideal op-amps have no input currents). With no current through R2 there is no voltage drop across it, and therefore the contribution to the output spectral noise density from this resistor is just:

$$v_{R_2} = e_{t2} = \sqrt{4kT\,R_2} \quad \mathrm{V}/\sqrt{\mathrm{Hz}} \qquad (31.26)$$

Note that while R_2 is often a larger resistor than R_1, the gain of the op-amp stage has no effect on the noise contribution of R_2. As a result it's usually the input resistance R_1, rather than the feedback resistance R_2 which is responsible for more of the output noise in this circuit.

Then there are the three noise sources associated with the op-amp itself to consider. Firstly, consider the noise current on the non-inverting input:

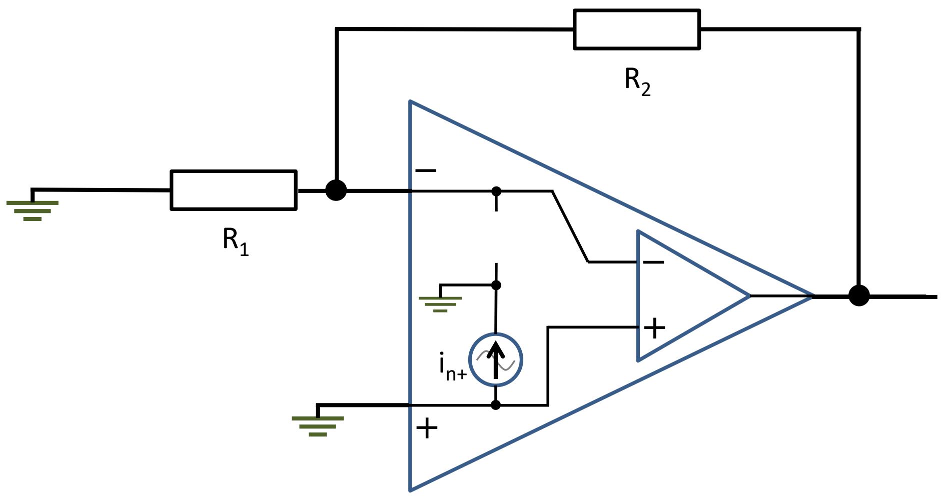

Figure 31.15 Showing the contribution of the first noise current to the noise of an inverting amplifier

Here the current source has no effect on the voltage on the inverting input, since it is tied directly to ground, so this noise source does not affect the voltages at either op-amp input, and hence makes no contribution to the total output noise. (Note that there is a trade-off here: connecting the non-inverting input directly to ground does minimise the noise, but doesn't eliminate the effect of the input bias current, so it results in an increased offset voltage on the output of the op-amp. Which effect is more important depends on the application.)

For the other noise current (the one on the inverting input):

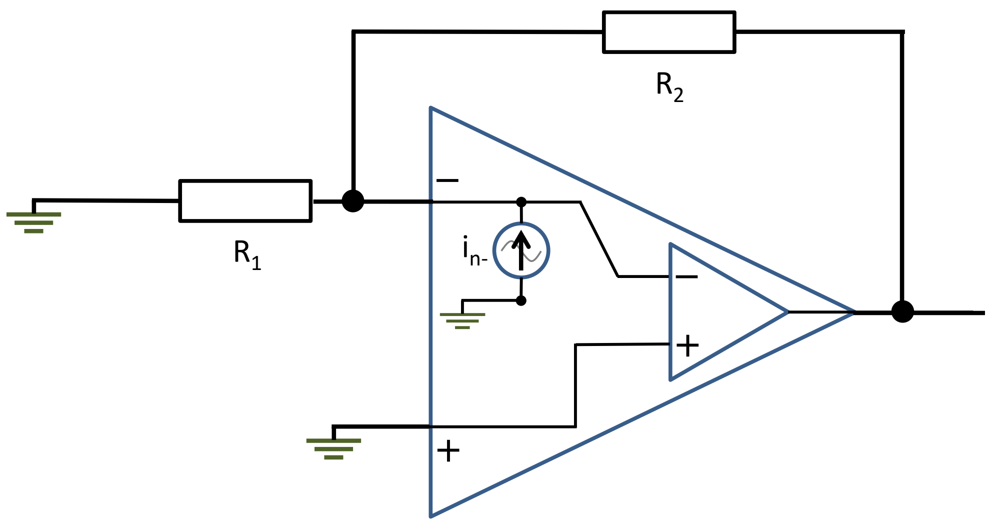

Figure 31.16 Showing the contribution of the second noise current to the noise of an inverting amplifier

the effect on the output can be determined to be:

$$v_{i_n} = i_n R_2 \quad \mathrm{V}/\sqrt{\mathrm{Hz}} \tag{31.27}$$

(The input resistor R_1 has no effect here since both ends of this resistor are at ground, so there is no current flowing through it. All the current in must therefore flow through the feedback resistor R_2, resulting in an output voltage of i_nR_2).

Finally, the effect of the op-amp's noise voltage:

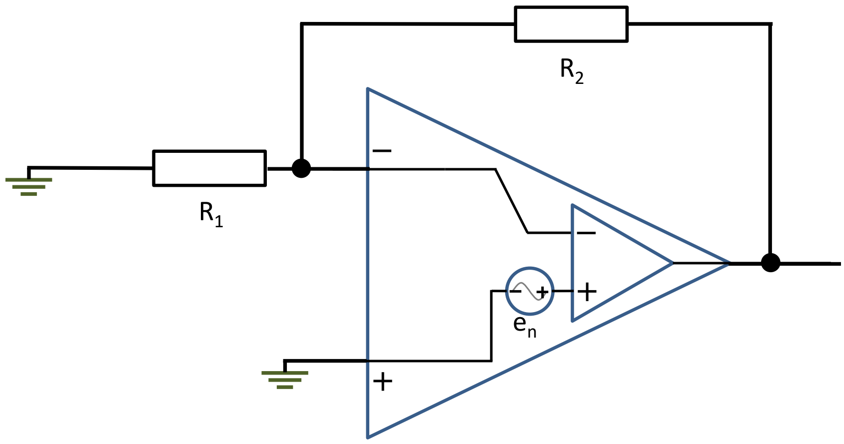

Figure 31.17 Showing the contribution of the op-amp noise voltage on the noise of an inverting amplifier

and this circuit has the form of a non-inverting op-amp, with the output given by:

$$v_{en} = e_n\left(1+\frac{R_2}{R_1}\right) \ \mathrm{V}/\sqrt{\mathrm{Hz}} \tag{31.28}$$

To work out the total noise voltage spectral density at the output, all these noise sources have to be squared to determine the noise power, added, then the square root taken to produce an expression for the total noise voltage:

$$\begin{aligned} n_{rms} &= \sqrt{v_{R_1}{}^2 + v_{R_2}{}^2 + v_{i_n}{}^2 + v_{en}{}^2} \\ &= \sqrt{R_2{}^2\frac{4kT}{R_1} + 4kTR_2 + i_n{}^2R_2{}^2 + e_n{}^2\left(1+\frac{R_2}{R_1}\right)^2} \ \ \mathrm{V}/\sqrt{\mathrm{Hz}} \end{aligned} \tag{31.29}$$

Putting in some typical numbers for a TL071 inverting amplifier with an input resistor of 10k and a feedback resistor of 100k (which produces an amplifier with a gain of -10), gives a predicted output noise spectral density of (at 290K):

$$n_{rms} = \sqrt{1.6\times10^{-14} + 1.6\times10^{-15} + 1\times10^{-18} + 3.92\times10^{-14}} \ \ \mathrm{V}/\sqrt{\mathrm{Hz}} \tag{31.30}$$

Note that the two dominant contributions to the noise are the 10k input resistor, and the amplifier's noise voltage. The other contributions are negligible in comparison. The total noise spectral density at the output would be expected to be:

$$\begin{aligned} n_{rms} &= \sqrt{1.6\times10^{-14} + 1.6\times10^{-15} + 1\times10^{-18} + 3.92\times10^{-14}} \\ &= 238 \ \ \mathrm{nV}/\sqrt{\mathrm{Hz}} \end{aligned} \tag{31.31}$$

so over an audio frequency range of 10 Hz to 18 kHz, this would predict a total noise rms voltage of:

$$n_{rms} = 238\ \frac{\text{nV}}{\sqrt{\text{Hz}}} \times \sqrt{17990\ \text{Hz}} = 31.9\ \mu\text{V} \tag{31.32}$$

Note that the second-largest contribution to this noise is the noise associated with the input resistance R_1. If both R_1 and R_2 were decreased by a factor of 10, the noise would be reduced but the gain would stay the same. So why not do this?

The problem is the output impedance of the previous stage. R_1 sets the input impedance of this stage, and if it is too low, then it will lower the input voltage due to the high currents causing too much voltage drop over the Thévenin equivalent output impedance of the previous stage. The noise will be reduced, but so will the signal level. Ultimately, it's the signal-to-noise ratio that is important, not the raw noise level[12], and decreasing the input resistor by too much can make the signal-to-noise ratio worse.

31.8.2 The effect of 1/f noise

At least, this works if the noise spectral density is white. Real op-amps, however, suffer from higher noise voltages at low frequencies, so the analysis has to be done slightly differently.

For the TL071, the 1/f noise is expressed in the datasheet as an equivalent input noise voltage over the frequency range from 10 Hz to 10 kHz, so for the noise contribution from 10 kHz to 18 kHz equation (31.31) works fine, and we can predict a total contribution to the noise from these frequencies of:

$$n_{rms\_HF} = 238\ \frac{\text{nV}}{\sqrt{\text{Hz}}} \times \sqrt{8000\ \text{Hz}} = 21.3\ \mu\text{V} \tag{31.33}$$

For the frequencies from 10 Hz to 10 kHz however, we need to return to equation (31.29) and now include the bandwidth, to give the total noise rms voltage:

$$n_{rms\_LF} = \sqrt{R_2^{\,2}\frac{4kT}{R_1}B + 4kTR_2B + i_n^{\,2}R_2^{\,2}B + e_n^{\,2}\left(1+\frac{R_2}{R_1}\right)^2 B}\ \ \text{V} \tag{31.34}$$

For the noise contribution from 10 Hz to 10 kHz, we can replace the contribution due to the input noise voltage with one in terms of the total input noise voltage over this bandwidth:

$$n_{rms\_LF} = \sqrt{R_2^{\,2}\frac{4kT}{R_1}B + 4kTR_2B + i_n^{\,2}R_2^{\,2}B + e_{n10\_10k}^{\,2}\left(1+\frac{R_2}{R_1}\right)^2}\ \ \text{V} \tag{31.35}$$

where e_{n10_10k} is the input noise voltage over this 10 Hz to 10 kHz bandwidth. Putting in the numbers here then gives:

---

[12] There is an optimum level for the input impedance in circuits like this; if you have a bit of spare time you might like to calculate what it is.

$$n_{rms\_LF} = \sqrt{1.6\times10^{-10} + 1.6\times10^{-11} + 1.0\times10^{-14} + 1.9\times10^{-9}} \ \ \text{V}$$
$$= 45.6 \ \ \mu\text{V} \tag{31.36}$$

and it becomes clear that the dominant source of noise in the output is the 1/f noise over this frequency range.

The total noise in the bandwidth from 10 Hz to 18 kHz is then:

$$n_{rms} = \sqrt{21.3^2 + 45.6^2} = 50.3 \ \ \mu\text{V} \tag{31.37}$$

This is significantly larger than the 31.9 uV calculated without including the effects of the 1/f noise; in fact the 1/f noise is the single largest factor in the output noise for this amplifier.

This leaves the question "what about the noise at frequencies below 10 Hz"? In theory at least the phenomenon of 1/f noise would predict an infinite noise spectral density at 0 Hz, and hence an infinite amount of noise, which in practice doesn't happen. However, there can be large amounts of noise at very low frequencies, and circuit techniques to remove this noise are required if high accuracy is needed at very low frequencies[13].

31.8.3 Analysing an op-amp differential amplifier stage for noise

Of most interest in the labs is the performance of a differential input amplifier circuit. This follows the same method, it's just a little more tedious since there are now seven noise sources to consider:

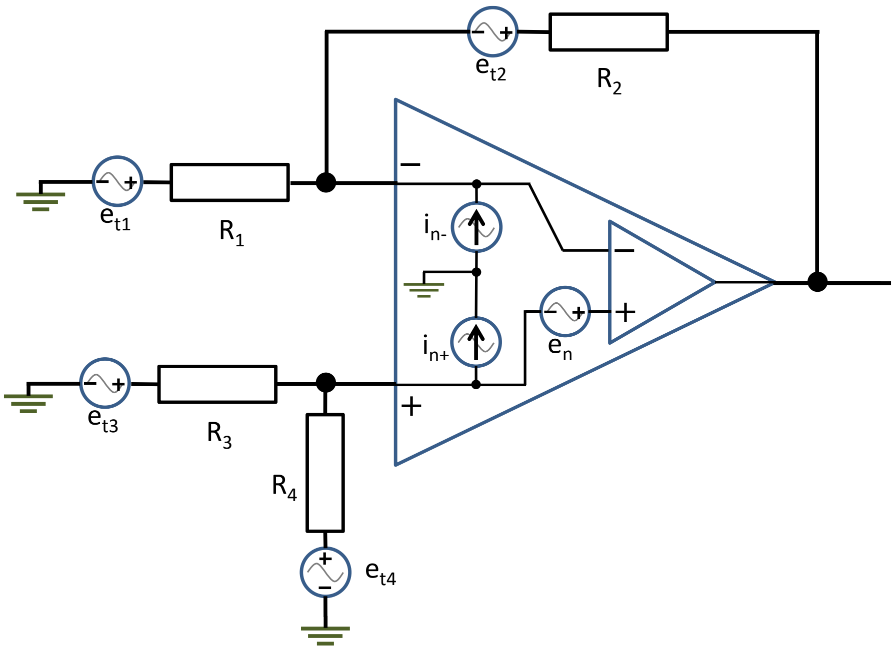

Figure 31.18 The noise sources in a differential-input op-amp amplifier

[13] For example, the use of "chopper-stabilised" op-amps which act to cancel out the low-frequency noise at their outputs at the expense of smaller bandwidths.

For cases where the gain of the stage is at least ten, and for op-amps with very low input currents (such as the TL071 or MCP6291), four of these sources can be neglected (i_{n+}, i_{n-}, e_{t2} and e_{t4}), and the remaining noise calculated to be:

$$n_{rms} = \sqrt{v_{R_1}{}^2 + v_{R_3}{}^2 + v_{en}{}^2}$$
$$= \sqrt{R_2{}^2 \frac{4kT}{R_1} B + 4kTR_2 B + e_n{}^2 \left(1 + \frac{R_2}{R_1}\right)^2 \left(B - 10^4\right) + e_{1/f}{}^2 \left(1 + \frac{R_2}{R_1}\right)^2} \quad \text{V} \tag{31.38}$$

Note that to save a bit of time here, I've separated out the contributions from the equivalent input noise up to 10 kHz (where the 1/f figure should be used) and over 10 kHz (where the figure in nV/√Hz should be used).

Putting in the numbers for a circuit with $R_1 = R_3 = 10k$, $R_2 = R_4 = 360k$, and using the parameters for an MCP6291 op-amp:

Noise						
Input Noise Voltage	E_{ni}	—	4.2	—	μV_{P-P}	f = 0.1 Hz to 10 Hz
Input Noise Voltage Density	e_{ni}	—	8.7	—	nV/√Hz	f = 10 kHz
Input Noise Current Density	i_{ni}	—	3	—	fA/√Hz	f = 1 kHz

Figure 31.19 The noise parameters from the MCP6291 datasheet

$$n_{rms} = \sqrt{2.9\times10^{-10} + 2.9\times10^{-11} + 4.9\times10^{-11} + 1.4\times10^{-9}} \quad \text{V}$$
$$= 42 \ \ \mu\text{V} \tag{31.39}$$

and note that again, in this case, it's the 1/f noise which is the largest single contribution to the noise over the audio frequency range.

31.8.4 A single equivalent input noise figure for amplifiers

Op-amps are specified with both input voltage and input current noise sources to allow the designer to predict the output noise for a range of amplifier configurations and gains. However, once the amplifier has been designed and the total output noise calculated, this can be converted back into a single equivalent input noise source.

All that is required is to take the total output noise voltage and divide by the gain provided by the stage to its input signal. Once this is done, the amplifier can be modelled as a perfect, noiseless amplifier, with an input given by the sum of the actual input signal and the equivalent input noise.

For example, in the inverting-amplifier case considered above where the op-amp had a total output noise of 50.3 µV, a feedback resistor of 100k and an input resistor of 10k, this can be modelled in terms of noise as an amplifier with a power gain of 100 (-10^2) and an equivalent input noise power of 25.3 pV2 (so that after this noise power is amplified by a factor of 100 it becomes 2.53nV^2 = 50.3 µV^2.)

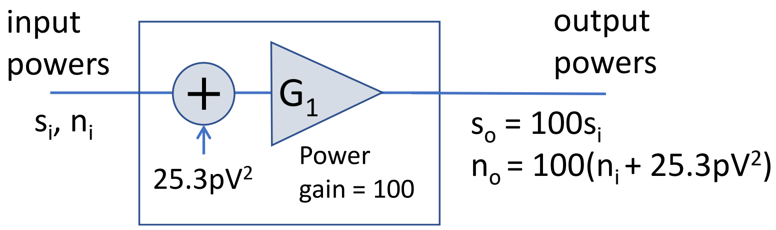

Figure 31.20 The noise parameters from the MCP6291 datasheet

This technique allows the output noise to be calculated easily given the noise added by the amplifier, and the noise power in the input signal. Note that it's the powers and power gains which are used here, the noises contribution of the amplifier is being added to any noise already present in the input signal, and when adding independent noise sources it's the powers that have to be added, not the amplitudes.

31.9 Multiple stages of noise

The final part of this short note deals with the issue of combining noise stages. In many cases (for example, a multiple-stage audio amplifier with the gain spread over the stages to increase the bandwidth), the total noise of a system is the sum of the contributions of several amplifier stages.

The easiest way to calculate how much noise appears at the output of a chain of noise-generating sub-systems is to consider the noise added by each stage, then move the equivalent input noise back to the beginning of the entire chain, in such a way that their contribution to the eventual output noise level is kept the same. Note that it's the noise powers and power gains that are used here, since we'll be adding up the noise contributions of the amplifier stages, and to add independent noise sources you have to add the noise powers.

So, for example, a three-stage amplifier which has power gains of G_1, G_2 and G_3 and equivalent input noise sources n_{e1}, n_{e2} and n_{e3} at the three stages:

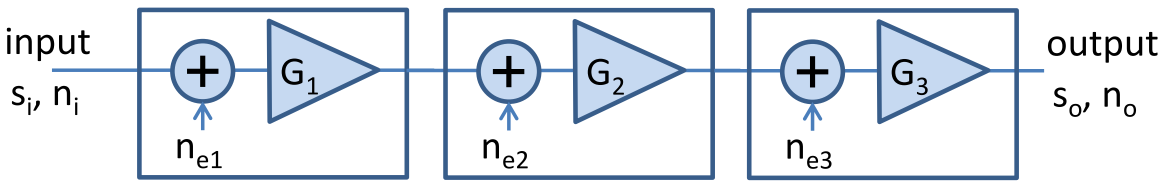

Figure 31.21 A three stage amplifier, with equivalent input noise at each stage

is equivalent (in terms of output noise and total gain) to:

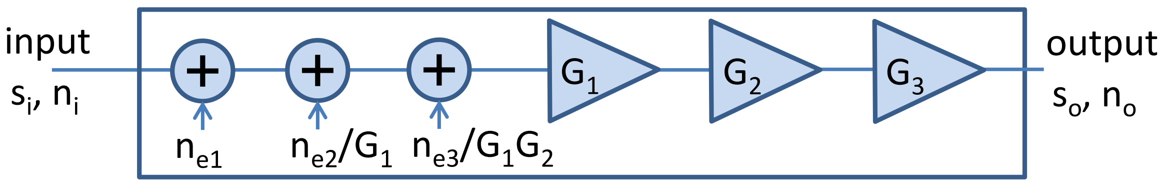

Figure 31.22 A three stage amplifier, with equivalent input noise referred to the input of the first stage

Note that to move the equivalent noise sources back to the start of the chain, they must be divided by the gains of the components they are moved through, so that the amount of noise generated at the output of the whole chain in the same. In the first case, we have:

$$n_o = \left(\left(\left(n_i + n_{e1}\right)G_1 + n_{e2}\right)G_2 + n_{e3}\right)G_3 = \left(n_i + n_{e1}\right)G_1G_2G_3 + n_{e2}G_2G_3 + n_{e3}G_3 \tag{31.40}$$

where n_i is the noise power in the input signal, and in the second (equivalent case):

$$n_o = \left(n_i + n_{e1} + \frac{n_{e2}}{G_1} + \frac{n_{e3}}{G_1G_2}\right)G_1G_2G_3 = \left(n_i + n_{e1}\right)G_1G_2G_3 + n_{e2}G_2G_3 + n_{e3}G_3 \tag{31.41}$$

which is clearly the same thing.

This allows us to define an equivalent input noise power for the entire three-amplifier system:

$$n_e = \left(n_{e1} + \frac{n_{e2}}{G_1} + \frac{n_{e3}}{G_1G_2}\right) \tag{31.42}$$

since this, when added to the input noise power and multiplied by the power gain of the system, gives the total output noise power.

(Again, note that we are adding noise powers here, because we assume that the noise contributions of the individual components are independent. n_e is a power, measured in Watts. G is a power gain, for example the square of the voltage gain.)

This is a form of a well-known formula known as the Friis formula[14]. It shows that as long as the gain of the first stage (G_1) is large enough, it's only the noise contribution of the first stage (n_{e1}) which makes a significant contribution to the overall system noise. This is why so much care is taken to minimise the noise of the first stage of multi-stage amplifiers (including, in some extreme cases, cooling them with liquid nitrogen).

31.10 Summary: the most important things to know

- Noise can be intrinsic, extrinsic or quantisation.
    - Intrinsic noise comes from the random movement of electrons in resistors (*Johnson* or *thermal noise*), or the currents in a diode (*shot noise*).
    - Extrinsic noise comes from variations on the power supply, communications cables or from external electric or magnetic fields.
    - Quantisation noise arises from the process of analogue-to-digital conversion
- Usually, intrinsic noise is white and Gaussian, extrinsic noise is neither white nor Gaussian, quantisation noise is white but not Gaussian.
    - "White" noise implies having the same power spectral density
    - "Gaussian" noise implies the probability density function of the noise is Gaussian
    - One exception is 1/f noise, which is not white but occurs in op-amp noise models
- When adding two independent noise sources, the squares of the individual rms noise voltages should be added to give the square of total rms noise.

[14] Named after the Danish engineer Harold Friis.

- Op-amps can be modelled as having three equivalent input noise sources: two currents (one per input) and one voltage.
- In a multiple-stage amplifier, most of the output noise is usually contributed by the first stage of the amplifier, so it is important to design this stage for minimum noise.