standalone first draft to be merged with noise analyzer document

1 Nyquist theorem and thermal noise power

Thermal noise, or Johnson-Nyquist noise, is the electrical noise generated by random thermal agitation of the charge carriers in an electrical conductor at equilibrium. The direction of the movement of the charge carriers is changed by collisions with, for example, the host crystal, impurities and other electrons, resulting in a random movement at zero voltage (Brownian motion), and a certain degree of randomness in their movement at finite voltages. This random movement gives rise to fluctuations in the net current in the conducting material. These fluctuations time-average to zero, but the time-average of the *square* of the fluctuations is not. If this non-zero variance of the voltage results from random collisions due to thermal agitation, one would anticipate that the distribution of voltage about its mean would be gaussian, which is what we aim to illustrate here.

Consider an ideal (lossless) and very long transmission line of length L which is terminated on each end by an identical resistance R. The entire system is being maintained in equilibrium at temperature T. We've constructed the transmission line so its characteristic impedance is also R, and everything is perfectly matched – any voltage wave propagating down the line is perfectly absorbed by the terminating resistor R without reflection. The resistor then is analogous to a one dimensional black body. Let a voltage wave of the form

$$V = V_o e^{i(kx-\omega t)} \tag{1}$$

propagate along the transmission line, with velocity $c' = \omega/k$. As one does in discussions of black body radiation, we'd now like to count the possible modes in the transmission line between the ends at $x = 0$ and $x = L$. We assume the boundary condition $V(x = L) = V(x = 0)$ on all propagating waves, which like the usual standing wave problem gives $kL = 2\pi n$ where n is an integer. Planck's hypothesis then tells us the mean energy in each mode at frequency f in a range $[f, f + df]$:

$$\epsilon(f) = \frac{hf}{e^{hf/k_BT} - 1} \Longrightarrow k_BT \quad \text{for } hf \ll k_BT \tag{2}$$

We will take the limit $hf \ll k_BT$ so that $\epsilon(f) = k_BT$. The *power* in each each mode is then energy per unit time, but the frequency range df implies a time interval $dt = 1/df$, meaning the power in each mode at frequency f in a bandwidth df can be written as $k_BT\,df$. The principle of detailed balance requires that each resistor is both emitting *and absorbing* this power if we are in thermal equilibrium.

If R_1 is emitting a power $k_BT\,df$, this power is being transmitted without loss down the transmission line and absorbed perfectly by R_2. From the point of R_2, this power can be considered as an instantaneous *voltage difference* applied from R_1, δv_{1f}. Since $R_1 = R_2$ and the two are in parallel, the current sourced through R_2 is $I_{1f} = \delta v_{1f}/2R_1$. The power generated in R_2 due to the radiated power from R_1 is then

$$I_{1f}^2 R_1 = R_1\left(\frac{\delta v_{1f}}{2R_1}\right)^2 = \frac{(\delta v_{1f})^2}{4R_1} \tag{3}$$

Since we require that the power emitted and absorbed by R_2 is the same, and since R_1 and R_2 are identical, the emitted power from R_1 and absorbed power from R_2 must be the same:

$$\frac{(\delta v_{1f})^2}{4R_1} = k_BT\,df \tag{4}$$

$$(\delta v_{1f})^2 = 4k_BTR_1\,df \tag{5}$$

This is the root mean square voltage, or variance of the voltage, produced due to thermal agitation at a temperature T in a resistor R_1 in a bandwidth of frequencies df wide. A more common way to view this is

$$\frac{(\delta v_{1f})^2}{df} \equiv S_v = 4k_B T R_1 \tag{6}$$

where S_v is the voltage *spectral density*, or the variance per unit frequency. It has units of signal (volts) squared per unit frequency. The variance of the voltage about its mean if our measurement encompassed frequencies from f_1 to f_2, with $\Delta f = f_2 - f_1$, would then be

$$\langle v^2 \rangle = \int_{f_1}^{f_2} S_v \, df = 4k_B T R_1 \Delta f \tag{7}$$

In short: the root mean square (rms) fluctuations of the voltage we measure due to random thermal motion is proportional to temperature, resistance, and the range of frequencies allowed in our measurement. That is: there is information in the "noise." As expected, for random thermal noise it is frequency-independent ("white noise") so long as we are in the classical limit $hf \ll k_B T$. What is interesting is that we have discovered a way to measure a resistance without passing a current: simply measure the voltage on a resistor at a known temperature as a function of time, and the rms voltage fluctuations will be proportional to resistance. This has applications in very low temperature thermometry. More importantly, by considering the problem a little more deeply we can use electrical noise to gain additional insight on the physics of electrical transport. And, of course, in order to build devices and instruments as immune to noise as possible, we need to understand the physics of noise better.

2 Signal power and spectral density

Consider a time-domain signal $x(t)$. We can also describe this signal in terms of its frequency components by considering its Fourier transform $x(f)$ and the inverse transform:

$$x(f) \equiv \mathcal{F}\{x(t)\} = \int_{-\infty}^{\infty} x(t) e^{-2\pi i f t} \, dt \tag{8}$$

$$x(t) = \mathcal{F}^{-1}\{x(f)\} = \int_{-\infty}^{\infty} x(f) e^{2\pi i f t} \, df \tag{9}$$

If you're note familiar with the idea of switching between the frequency and time domain, you'll pick it up soon enough. The time average power of the signal $x(t)$ centered around some arbitrary time t_0 is

$$P = \lim_{T\to\infty} \frac{1}{T} \int_{t_0-T/2}^{t_0+T/2} |x(t)|^2 \, dt \tag{10}$$

For the sake of convenience, let us define a function $x_T(t)$ which is equal to $x(t)$ in $t \in [-T/2, T/2]$ and zero elsewhere. That allows us to perform the integral over all times:

$$P = \lim_{T\to\infty} \frac{1}{T} \int_{-\infty}^{\infty} |x_T(t)|^2 \, dt \tag{11}$$

Parseval's theorem states that the power in a signal $x_T(t)$ integrated over all times is the same as the power of its Fourier transform $x_T(f) = \mathcal{F}\{x_T(t)\}$ over all times. That is:

$$P = \lim_{T\to\infty} \frac{1}{T} \int_{-\infty}^{\infty} |x_T(t)|^2 \, dt = \lim_{T\to\infty} \frac{1}{T} \int_{-\infty}^{\infty} |x_T(f)|^2 \, df \tag{12}$$

This makes sense: whether we divide the signal in the frequency or time domain, either way we are *physically* talking about the power contained in the entire signal, which must be the same no matter the representation. Now we define

$$S_{xx}(f) = \lim_{T\to\infty} \frac{1}{T} |x_T(f)|^2 \qquad \Longrightarrow \qquad P = \int_{-\infty}^{\infty} S_{xx}(f) \, df \tag{13}$$

Thus S_{xx} must be the power per unit frequency - the spectral density - just as we found for thermal noise.

3 Correlations and convolutions

Now consider the squared magnitude of the frequency-domain signal, $|x_T(f)|^2$.

$$|x_T(f)|^2 = x_T^*(f) x_T(f) \tag{14}$$

here * indicates complex conjugation. An interesting thing results if we make use of the *cross correlation theorem*, which is analogous to the convolution theorem. First, the cross correlation of continuous functions f and g is defined as

$$R_{ff}(t) = (f \star g)(t) \equiv \int_{-\infty}^{\infty} f^*(\tau) g(t+\tau) \, dt \tag{15}$$

For signals that are not strictly periodic, we can define the cross correlation over a period of time as above,

$$R_{ff}(t) = (f \star g)(t) \equiv \lim_{T\to\infty} \int_{0}^{\infty} f^*(\tau) g(t+\tau) \, dt \tag{16}$$

where f^* means the complex conjugate of f. If we consider f and g as time domain signals, the cross-correlation is a measure of similarity of the two signals as a function of their displacement or lag. It is similar to a convolution, and is often used to search a longer signal for a shorter pattern. The cross correlation of a signal with itself, $(f \star f)(t)$ is the *auto-correlation* function, and it reveals periodicities in the signal.

The cross-correlation obeys a relationship analogous to the convolution:

$$\mathcal{F}\{f \star g\} = \Big(\mathcal{F}\{f\}\Big)^* \mathcal{F}\{g\} \qquad \text{noting} \quad \Big(\mathcal{F}\{f(t)\}\Big)^* = \mathcal{F}\{f^*(-t)\} \tag{17}$$

Then going back to Eq. 14,

$$|x_T(f)|^2 = x_T^*(f)x_T(f) = \mathcal{F}\{x_T^*(t)\}\mathcal{F}\{x_T(t)\} = \Big(\mathcal{F}\{x_T(-t)\}\Big)^*\mathcal{F}\{x_T(t)\} = \mathcal{F}\{(x_T \star x_T)(t)\} \tag{18}$$

Applying this to Eq. 13, and using the definition of $\mathcal{F}$,

$$S_{xx}(f) = \lim_{T\to\infty}\frac{1}{T}|x_T(f)|^2 = \lim_{T\to\infty}\frac{1}{T}\mathcal{F}\{(x_T \star x_T)(t)\} = \lim_{T\to\infty}\frac{1}{T}\int_{-\infty}^{\infty}\Big[(x_T \star x_T)(t)\Big]e^{-2\pi i f\tau}\,d\tau \tag{19}$$

$$S_{xx}(f) = \int_{-\infty}^{\infty}\Big[\lim_{T\to\infty}\frac{1}{T}(x_T \star x_T)(t)\Big]e^{-2\pi i f\tau}\,d\tau = \int_{-\infty}^{\infty}\Big[\lim_{T\to\infty}\frac{1}{T}\int_{-\infty}^{\infty}x_T^*(\tau)x_T(t+\tau)\,dt\Big]e^{-2\pi i f\tau}\,d\tau \tag{20}$$

$$S_{xx}(f) = \int_{-\infty}^{\infty}R_{xx}(\tau)e^{2\pi i f\tau}\,d\tau = \mathcal{F}\{R_{xx}(\tau)\} = R_{xx}(f) \tag{21}$$

This is the result we wanted: for a time-domain signal $x(t)$ in some time range T, the noise spectral density is the Fourier transform of the autocorrelation of the signal. The noise spectral density $S_xx(f)$ is the power of fluctuations per unit frequency in a frequency bandwidth $[f, f+df]$, so it is telling us about the energy-dependence of fluctuations in our system. The auto-correlation of the signal is nothing more than the mathematical manipulation of a voltage versus time trace for our signal. Thus, a simple time-domain measurement and some post-processing allows a measurement of the fluctuation frequency spectrum.

Now there are details. Our measurement of $x(t)$ will have in reality two characteristic times: the total time for a given measurement, T, and the sampling time between individual measurement points. The former gives a lower-frequency limit to our knowledge of $S_{xx}(f)$ of $f_L = 1/T$. For example, a 10 s duration measurement gives us a measure of $S_{xx}(f)$ down to 0.1 Hz. On the other end, if the signal is discretely sampled $2B$ times per second, such that measurement points are $1/2B$ seconds apart, the Shannon-Nyquist theorem tells us that B Hz is the highest frequency that can be reconstructed. For example, if we measure 1000 samples per second (1 ms time intervals), we can reconstruct information up to 500 Hz.

4 Eliminating external noise: cross correlation spectra

Of course, a significant practical problem arises when attempting to measure noise: how do you ensure you are measuring the noise of the device of interest and *not* your wires, amplifiers, meters, etc? As it turns out, there are some rather clever signal processing techniques that allow you to isolate the noise from the device of interest, independent of the hookup wires, amplifiers, etc., to an almost arbitrary degree. The price is a computational cost and additional hardware.

Essentially, we measure the voltage versus time on the device of interest *twice*, independently. From these two time domain signals, we can compute their *cross correlation* and its Fourier transformation to find the noise spectral density (noise power per unit frequency versus frequency). If we integrate for a sufficiently long time, only the frequency components that the two measurements share will survive in the cross-correlation signal. Physically, the only frequency components that the two measurements will have in common are those arising from their common signal path, which you can restrict to only the sample under study relatively easily.

First, the cross-correlation spectrum. Consider two continuous time domain signals $x(t)$ and $y(t)$. We can define their cross correlation function as:

$$R_{xy}(t) = \big(x \star y\big)(t) \equiv \lim_{T\to\infty}\int_{-T/2}^{T/2}x(\tau)y^*(\tau - t)d\tau \tag{22}$$

where * denotes complex conjugation. According to the Wiener-Khinchin theorem, sketched above, the power spectral density (PSD), $S_{xy}(f)$, is the Fourier transform ($\mathcal{F}$) of the cross correlation function: ?

$$S_{xy}(f) = \mathcal{F}\left[R_{xy}(t)\right] \tag{23}$$

The cross-correlation theorem states that $\mathcal{F}\left[f \star g\right] = \left(\mathcal{F}\left(f\right)\right)^* \mathcal{F}\left(g\right)$, which gives

$$S_{xy}(f) = \mathcal{F}\left[\left(x \star y\right)(t)\right] = \left(\mathcal{F}\left[x(t)\right]\right)^* \mathcal{F}\left[y(t)\right] \tag{24}$$

As a matter of practical convenience, we note that if $x(t)$ and $y(t)$ have Fourier transforms $\tilde{x}(f)$ and $\tilde{y}(f)$,

$$S_{xy}(f) = \mathcal{F}\left[x(t)\right]^* \mathcal{F}\left[y(t)\right] = \tilde{x}^*(f)\tilde{y}(f) \tag{25}$$

In the case of single-sided spectra (i.e., positive frequencies only), one must double the Fourier transform of both $x(t)$ and $y(t)$, giving:

$$S_{xy}(\omega) = 4\tilde{x}^*(\omega)\tilde{y}(\omega) \tag{26}$$

The question is, how does this help get rid of signal components that are not common to x and y? Let's compute a simple cross-correlation and we will see.

4.1 Auto correlation of a sinusoid

Let $x(t) = \sin t$. Then

$$R_{xx}(\tau) = \int_{-\infty}^{\infty} x(t-\tau)x(t)\,dt = \int_{-\infty}^{\infty} \sin\left(t-\tau\right)\sin t\,dt = \int_{-\infty}^{\infty} \left(\sin t\cos\tau - \sin\tau\cos t\right)\sin t\,dt \tag{27}$$

The second term in parenthesis has the form $\sin t \cos t$ and will integrate to zero over any integer number of periods.

$$R_{xx}(\tau) = \int_{-\infty}^{\infty} \sin^2 t\cos\tau\,dt = \frac{1}{2}\cos\tau \tag{28}$$

The autocorrelation of a sine wave as a function of the lag between it and a shifted copy of itself is proportional to the cosine of the lag, which is just telling us that the correlation is maximum when the signal and shifted copy are in phase with each other.

4.2 Cross correlation of two sinusoids

Let $x(t) = \sin\left(\omega_1 t\right)$ and $y(t) = \sin\left(\omega_2 t\right)$. You can convince yourself that since the cross-correlation already involves a lag that adding a phase to one of the signals changes nothing much. Then

$$R_{xy}(\tau) = \int_{-\infty}^{\infty} x(t-\tau)y(t)\,dt = \int_{-\infty}^{\infty} \sin\left(\omega_1 t - \omega_1\tau\right)\sin\left(\omega_2 t\right)dt \tag{29}$$

$$= \int_{-\infty}^{\infty} \left(\sin\left(\omega_1 t\right)\cos\left(\omega_1\tau\right) - \sin\left(\omega_1\tau\right)\cos\left(\omega_1 t\right)\right)\sin\left(\omega_2 t\right)dt \tag{30}$$

Again second term in parenthesis has the form $\sin t \cos t$ and will integrate to zero over any integer number of periods.

$$R_{xy}(\tau) = \int_{-\infty}^{\infty} \sin(\omega_1 t)\cos(\omega_1 \tau)\sin(\omega_2 t)\, dt = \cos(\omega_1 \tau)\int_{-\infty}^{\infty} \sin(\omega_1 t)\sin(\omega_2 t)\, dt \tag{31}$$

$$= \cos(\omega_1 \tau)\frac{1}{2}\delta(\omega_1 - \omega_2) \tag{32}$$

The orthogonality of the sin function means the integral is zero unless $\omega_1 = \omega_2$, in which case it is 1/2. Given that we can build any periodic function out of sinusoids, this implies an important result: *only the frequency components two signals have in common will give a non-zero cross-correlation.* In other words, the cross-correlation is a tool for picking out the shared frequency components in two signals. Now, this requires integrating for infinite time, which is hardly practical, but for "sufficiently long" integration times we can reduce the contribution from extraneous (non-shared) frequency components to be arbitrarily small.

What that means is that we can reject external noise from, say, amplifiers and wires and the environment and pick out only the noise on our device of interest, simply by measuring the signal twice independently.

4.3 Cross-correlation of model signal plus noise

Now imagine each of the signals has a common portion at frequency ω_c and each has an incoherent contribution not shared by the other channel at a different frequency ω_i. Keep in mind that the ω_c and ω_i are stand-ins for a Fourier series we could use to build any function - if the result works for one sinusoid, since our operations are all linear it will work for a superposition of them. Specifically, consider now

$$x(t) = A_{c1}\sin(\omega_c t) + A_1 \sin(\omega_1 t) \tag{33}$$

$$y(t) = A_{c2}\sin(\omega_c t + \delta) + A_2 \sin(\omega_2 t) \tag{34}$$

We will assume that the two measurements may have slightly different contacts such that a phase difference in the common signal can be introduced. Of course we could introduce phase shifts on the incoherent portions as well, but you can convince yourself it won't make a significant difference in the argument. The cross correlation is then

$$R_{xy}(\tau) = \int_{-\infty}^{\infty} \Big(A_{c1}\sin(\omega_c t - \omega_c \tau) + A_1 \sin(\omega_1 t - \omega_1 \tau)\Big)\Big(A_{c2}\sin(\omega_c t + \delta) + A_2 \sin(\omega_2 t)\Big)\, dt \tag{35}$$

At some point I will type up the details, but expand the above into four terms. Then remember that τ is not an integration variable! All terms which involve the product of sinusoids at different frequencies like $\sin(\omega_1 t)\sin(\omega_c t)$ will integrate to zero over an integer number of cycles, so those terms can be neglected. The only surviving terms are those those with products of sin's of t at the same frequency, and only the first term in the expansion has this. The result is

$$R_{xy}(\tau) = \frac{1}{2}A_{c1}A_{c2}\cos(\omega_c \tau)\cos\delta \tag{36}$$

Provided we integrate for an infinite amount of time, only the common signal survives the cross correlation, and its magnitude scales with the phase shift between the two signals (so minimize that in your setup by placing the signal contacts close to each other). This means that S_{xy} will also only contain contributions from the common signal.

In short: measuring your signal twice and using the cross-correlation technique allows you to determine the noise spectral density due only to the shared signal path. Practically speaking, this means you can measure the noise of your device alone, while rejecting external noise to a degree mainly limited by your patience.