RSI Technical Note 029

Noise of the MicroSquid Dissolved Oxygen Measurement System

Rolf Lueck

2013-11-06

Rockland Scientific International Inc.
520 Dupplin Rd
Victoria, BC, CANADA, V8Z 1C1
www.rocklandscientific.com

Contents

1 Introduction **2**

2 Noise Signal – No Sensor Installed **3**

3 Noise Signal – Sensor Installed **5**

4 Noise Spectra **6**

5 Cumulative Noise Variance **8**

6 Noise In terms of oxygen Concentration **10**

7 Discussion **12**
- 7.1 Time- versus frequency-domain estimation . . . 12
- 7.2 Why is noise important to co-variance estimation? . . . 13

List of Figures

1 Noise time series – no probe installed . . . 3
2 Noise time series – no probe installed, close up. . . . 4
3 Noise time series – probe installed . . . 5
4 Noise Current Spectra . . . 6
5 Cumulative Noise Current Spectra . . . 8
6 Noise signal in terms of concentration . . . 10
7 Concentration Noise Spectra . . . 11
8 Cumulative Concentration Noise Spectra . . . 11
9 1024-second histograms of standard deviation . . . 12
10 128-second histograms of standard deviation . . . 13

1 Introduction

This document describes the measurement noise of the MicroSquid dissolved oxygen instrument (P055), both with and without an AMT oxygen sensor installed, and in terms of spectra and time-domain fluctuations.

For the first measurement, we collected data with no probe installed. This provides the noise that is attributable to the electronics alone. It is the best that can be achieved with the current version of the electronics (P055R00) regardless of the sensors that could be installed.

For the second test we installed the AMT (galvanic) sensor, serial number SN010. The sensor is completely immersed in nitrogen gas. A coupler is slipped over the probe guard and this coupler is connected to a Tygon tube. The tube is connected to a bottle of nitrogen gas with a slight over pressure (with respect to atmospheric pressure) so that the gas flows slowly and can escape through a vent hole in the coupling. This flow purges all oxygen from the sensor exterior. I am assuming that the added noise from the sensor is independent of the flow of oxygen into the sensor. I have not, so far, been able to conceive of a noise test with the sensor in oxygenated water, wherein the oxygen is kept at a constant level and without fluctuations.

The MicroSquid electronics is inside the standard housing that is currently used for the dissolved oxygen measurement system. It has been modified by extending the shielding throughout the entire hole in the front bulkhead. The probe holder, the extra shielding in the bulkhead, and the shield lining the inner wall of the pressure case are all connected to the analog ground of the isolated $\pm 15\,\mathrm{V}$ supply that energizes the sensor electronics. The entire instrument was wrapped in aluminum foil which was connected to the probe guard. The return of the exterior ($+15\,\mathrm{V}$) supply used to energize the instrument was also connected to the foil shield.

The output signal from the instrument was routed to our general purpose analog input board (GPIO) and sampled at a rate of $1024\,\mathrm{s}^{-1}$. The full-scale range of the input to the GPIO is 0 to 5.12V and the number of bits is 16. The data are deliberately over sampled to reduce the quantization (or sampling) noise so that the measured signal is dominated by noise from the instrument. The GPIO was mounted on a rail and was accompanied by a power supply board and a RSTRANS (remote serial transceiver) board that was used to communicate the data to a Notebook computer via a UTRANS (USB transceiver) board. The Notebook computer was energized by its battery charger and, therefore, not isolated from the 60 Hz mains of the laboratory. This was necessary because of the extended duration of the measurements. We did notice a very slight (≈ 0.3 counts) reduction in the standard deviation of the signal when we disconnected the battery charger from the Notebook computer.

2 Noise Signal – No Sensor Installed

The measured signal, with no sensor installed, is converted into equivalent input current noise because the instrument is designed to measure current from an oxygen sensor. The transformation of measured signal (in units of counts) into current is

$$I = \left(\frac{V_{FS}}{2^B} N - V_0 \right) \frac{1}{R/2} \tag{1}$$

where N is the sampled data from the analog-to-digital converter (a dimensionless integer), $V_{FS} = 5.120\,\mathrm{V}$ is the full-scale range of the analog-to-digital converter on the GPIO board, $B = 16$ is its number of bits, $V_0 = 0.125\,\mathrm{V}$ is the intensionally added offset of the output from the MicroSquid, $R = 5 \times 10^9\,\Omega$ is the value of the resistor used to transform current into a voltage, and the factor 2 is the amplitude reduction of the optical coupler between the isolated side of the sensor circuitry and the output from the instrument.

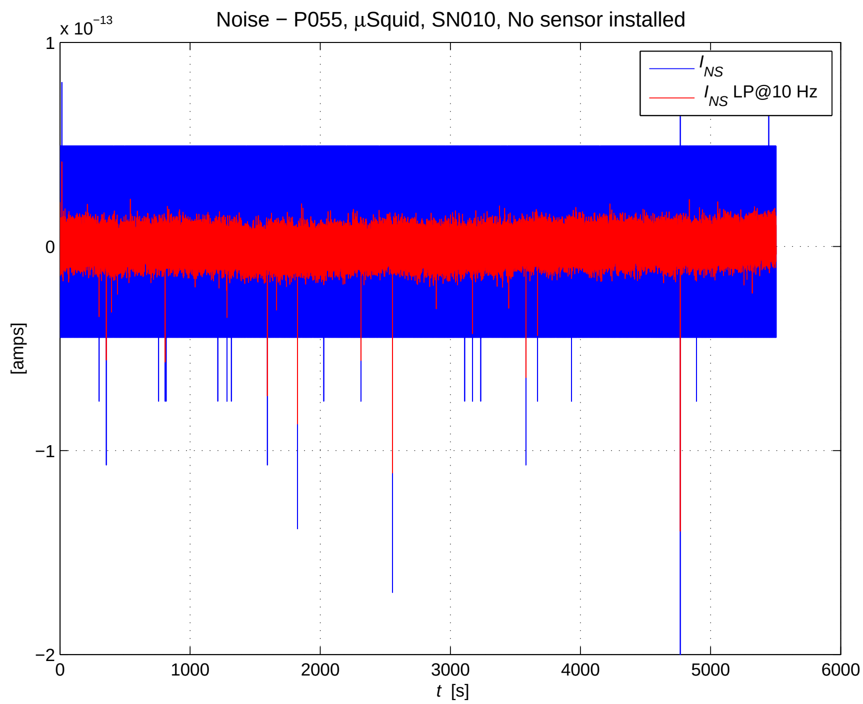

Figure 1: The time series of measured current when no probe is installed in the MicroSquid DO instrument. The blue trace is the full-bandwidth (512 Hz) signal. The red trace is the same signal after low-pass filtering at 10 Hz.

The full-bandwidth current signal is composed almost entirely of 2 values (blue trace, Figure 1) with a few extrema and mainly in the negative direction. The usable bandwidth of the instrument is 10 Hz and so I have also low-pass filtered the data at this frequency (red trace, Figure 1). The extrema are only partially reduced after filtering and this indicates that they are not isolated "fliers". A close up view of one of the extrema (Figure 2) shows that they last for about 0.1 seconds. This means that they will add noise to the spectrum at $\approx 10\,\mathrm{Hz}$ and higher. The reason for these extrema is unknown and such features were not detected later when a probe was installed.

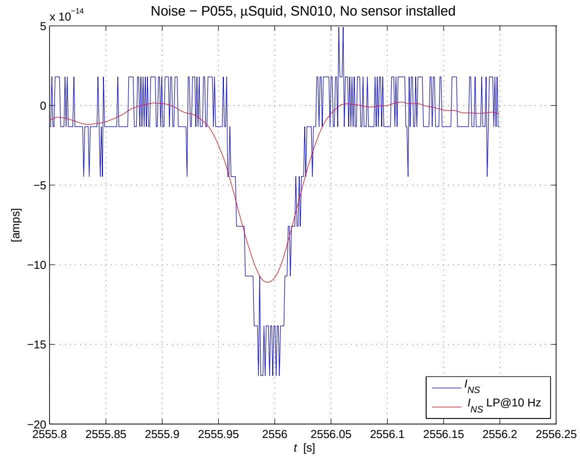

Figure 2: Same as Figure 1 but zoomed in on a single extrema.

3 Noise Signal – Sensor Installed

The measured signal, when a sensor is installed and is held in an atmosphere of nitrogen, shows a small drift over many hours ((blue trace, Figure 3)). A mathematical model of this

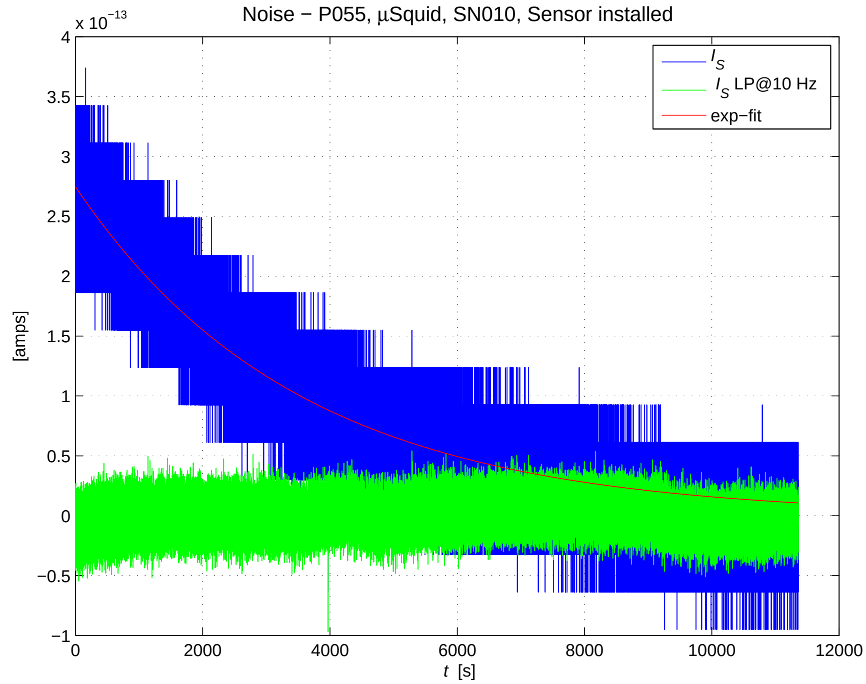

Figure 3: The time series of measured current when when an AMT probe is installed in the MicroSquid DO instrument. The blue trace is the full-bandwidth (512 Hz) signal. The red trace is a mathematical model of the drift. The green trace is the the blue trace minus the red trace after low-pass filtering at 10 Hz.

drift (red trace, Figure 3) is

$$I_{Drift} = A \exp\left(\frac{t}{\tau}\right) \tag{2}$$

where $A = 2.77 \times 10^{-13}$ A is the initial value and $\tau = 3500$ s is its relaxation time. I subtracted this mathematical model from the measured current, and then low-pass filtered it at 10 Hz to derive the zero-mean noise signal with a probe attached to the instrument (green trace, Figure 3). The peak-to-peak amplitude of the fluctuations is about twice that when there is no probe installed (red trace, Figure 1), if we ignore the fliers.

4 Noise Spectra

Spectra of the noise signal were computed using two methods. In order to resolve frequencies down to 1×10^{-3} Hz, I used segments that are 1024 s long, with 50 % overlap between adjacent segments. This provides the highest frequency resolution but reduced statistical certainty because the number of segments available for ensemble averaging is low – 27 and 12 for the measurements with and without a probe installed, respectively. These spectra are the thin traces in Figure 4 and are shown only for frequencies smaller than 0.1 Hz.

I also calculated the spectra using 128 s segments and these are the thick traces in Figure 4. Their statistical uncertainty is is reduced 8-fold compared to the thin traces. The theoretical

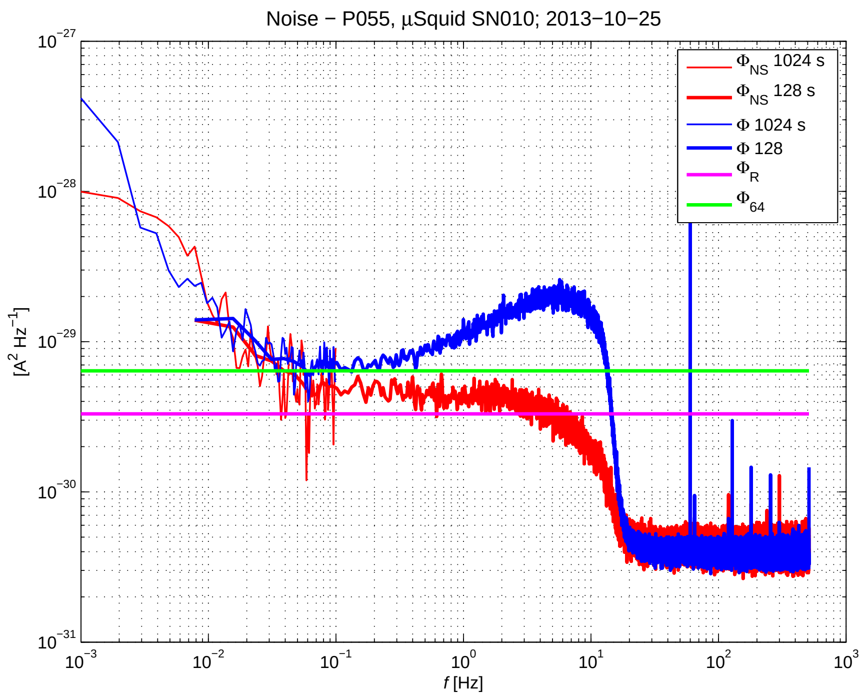

Figure 4: The spectra of the noise current with (blue) and without (red) the AMT sensor installed in the instrument. The thin (thick) lines are spectra based on 1024 s (128 s) segments, respectively. The noise of a $5 \times 10^9\,\Omega$ resistor (magenta) and a practical sampler running at 64 Hz (green).

limit of noise is imposed by thermodynamics. The resistor that is used to convert the current into a voltage itself produces noise because of the thermal agitation of the electrons within its crystal structure. This noise is usually call “thermal” or “Johnston” noise. The current noise spectrum of a perfect metal resistor is independent of frequency and has a value of

$$\Phi_R = \frac{4kT}{R} \;=\; 3.3 \times 10^{-30}\,\mathrm{A}^2\,\mathrm{Hz}^{-1} \tag{3}$$

where $k = 1.38 \times 10^{-23}\,\mathrm{m}^2\,\mathrm{kg}\,\mathrm{s}^{-2}\,\mathrm{K}^{-1}$ is the Boltzmann constant, $T = 295\,\mathrm{K}$ is the absolute temperature, and $R = 5 \times 10^9\,\Omega$ is the value of the resistor.

Without a probe installed, the noise spectrum of the instrument is just above the Johnson noise (magenta line in Figure 4) in the frequency range of 0.05 to 5Hz (Figure 4). For higher frequencies, the noise drops below the Johnson noise level because of the 10 Hz filter in the output of the instrument and because of intentional (2.7 pF) and stray (≈ 1 pF) capacitance at the input of the amplifier. The intentional capacitance is needed for circuit stability. For frequencies smaller than 0.05 Hz the noise spectrum rises – it is approximately proportional to $1/f$. This rise cannot be explained by either the current or the voltage noise of the integrated circuit used for the current-to-voltage conversion – it is not amplifier noise. It is likely (but still unproven) that this excess noise is due to the analog optical coupler used to isolate the sensor electronics from the output of the instrument.

When the AMT sensor is installed, the noise spectrum differs little from that without the sensor, for frequencies smaller than 0.1 Hz (blue traces, Figure 4). In the band from 0.1 to 10 Hz the noise is significantly ($\sim \times 10$) larger when a sensor is installed, compared to no sensor. For higher frequencies, the noise is the sampling noise of our analog-to-digital converter plus some small contributions from the 60 Hz mains in the laboratory and, for this range of frequencies, the noise differs little from that without a sensor installed.

It is worth comparing the noise spectra against the theoretical and practical noise imposed by a sampling system. An ideal analog-to-digital converter has a sampling noise spectrum of

$$\Phi_S = \frac{1}{12}\frac{\Delta^2}{f_N} \tag{4}$$

where Δ is the step size of the analog-to-digital converter, and f_N is the Nyquist frequency and equals half the sampling rate. The noise of an ideal sampler is independent of frequency. For the full-scale range and the number of bits used here, the sampling step size is 3.1×10^{-14} A. For the sampling rate used here, the ideal sampling noise spectrum is $1.6 \times 10^{-31}\,\mathrm{A}^2\,\mathrm{Hz}^{-1}$. Our sampler reaches $4 \times 10^{-31}\,\mathrm{A}^2\,\mathrm{Hz}^{-1}$ which can be deciphered by looking at the average spectral level at frequencies large compared to 10 Hz in Figure 4. The sampler in the Nortek Vector is only slightly inferior to ours. Thus, we can scale the sampling noise of Figure 4 up to where it would be if we had chosen to sample the data at a rate of $64\,\mathrm{s}^{-1}$ (the maximum rate available with the Vector). This noise level is $6 \times 10^{-30}\,\mathrm{A}^2\,\mathrm{Hz}^{-1}$ (green line in Figure 4). It is slightly above the Johnson noise and slightly below the noise when a sensor is installed.

5 Cumulative Noise Variance

The spectra of the previous section tell us how the variance of noise is distributed with respect to frequency. However, the spectra do not (directly) tell us what the variance is over a band of frequencies. The variance, $\overline{I^2}$, of a signal, $I(t)$, with a spectrum $\Phi_I(f)$ is

$$\overline{I^2} = \int_{f_0}^{f_1} \Phi_I \mathrm{d}f \tag{5}$$

for all frequencies between f_0 and f_1.

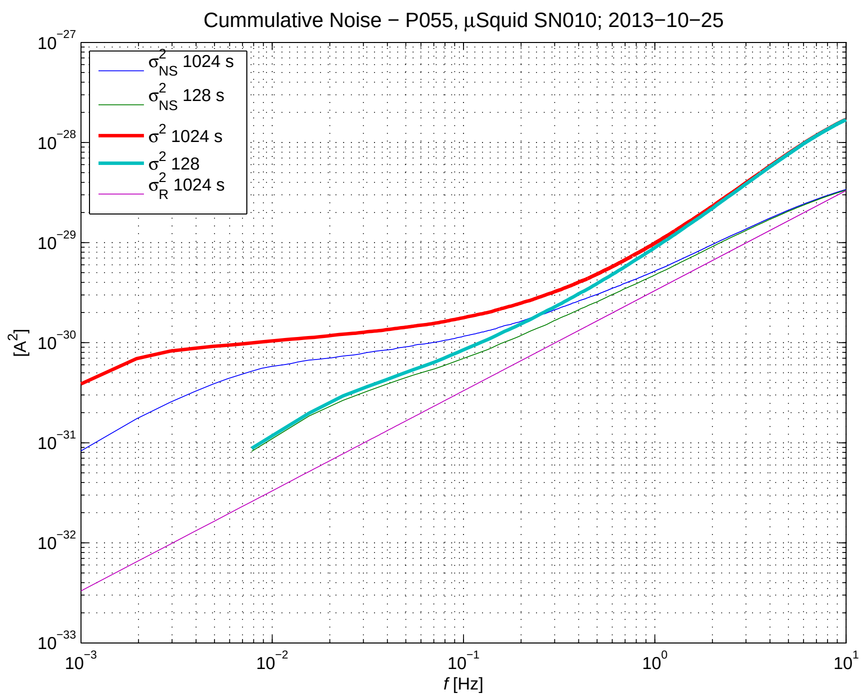

Figure 5: The cumulative noise, as a function of frequency, for spectral integration starting at 0.001 (thin blue and thick red) and 0.01 Hz (thin green and thick cyan). The subscript NS indicates noise without a sensor installed. The magenta line is the cumulative noise for an ideal resistor of $5 \times 10^9\,\Omega$.

I have computed the cumulative variance of the current noise by integrating the spectra from 0.001 to f, namely

$$\Psi_{0.001}(f) = \int_{0.001}^{f} \Phi_I(\eta)\mathrm{d}\eta \tag{6}$$

(Figure 5, thin blue and thick red curves), as well as from 0.01 to f, namely

$$\Psi_{0.01} = \int_{0.01}^{f} \Phi_I(\eta)\mathrm{d}\eta \tag{7}$$

(Figure 5, thin green and thick cyan curves), both without a sensor installed (which are identified by the subscript "NS") and with a sensor installed (which are not subscripted).

Integrating from 0.001 (or 0.01) to 10 Hz is equivalent to finding the mean-square of $I(t)$ over an interval of 1000 (or 100 s).

The duration of data used to make a single estimate of the vertical flux of oxygen is a matter of scientific choice but it usually ranges from 1 to several minutes. If, for example, a researcher did want to compute fluxes over intervals of 100 seconds, then the thick cyan curve in Figure 5 at 10 Hz tells us that the noise variance of each estimate is just under 2×10^{-28} A^2. The same result holds for fluxes spanning intervals of 1000 seconds (red curve), because the noise comes predominantly from frequencies higher than 0.1 Hz.

6 Noise In terms of oxygen Concentration

Expressing the measurement noise in terms of the fundamental quantity measured by the instrument, namely current, is of great interest to the engineering of an instrument. However, the user is more likely to want to know what the noise is in terms of the physical parameter that they wish to measure, namely the concentration of dissolved oxygen.

All of the above figures can be readily converted to oxygen concentration if we know the current produced by a sensor as a function of the concentration of oxygen in water. The relationship is quite linear and the current output is very small when the water is anoxic. That is, the current output is essentially proportional to the concentration of oxygen. I will use the typical current produced at 100 % concentration (the saturation concentration) at atmospheric pressure and room temperature for the conversion factor. A typical value for the AMT sensor is $I_{100} = 2000\,\mathrm{pA}$. To be clear, if I take the current signal (such as shown in Figures 1 – 3) and divide this signal by I_{100}, then this new signal is the fluctuations (or noise) relative to a concentration of 100 %. Therefore, this signal is dimensionless (Figure 6). It can be converted to any of the many other units used to quantify oxygen concentration (such as $\mathrm{mol\,L^{-1}}$) by multiplying this signal by value (in the other units) at 100 % concentration.

Figure 6: The signals without a sensor installed (red) and with the AMT sensor installed, in terms of the signal produced at 100 % water concentration.

Likewise, spectra and cumulative spectra (such as Figures 4 and 5) can be expressed in terms of concentration squared by dividing these current spectra by I_{100}^2 (Figures 7 and 8).

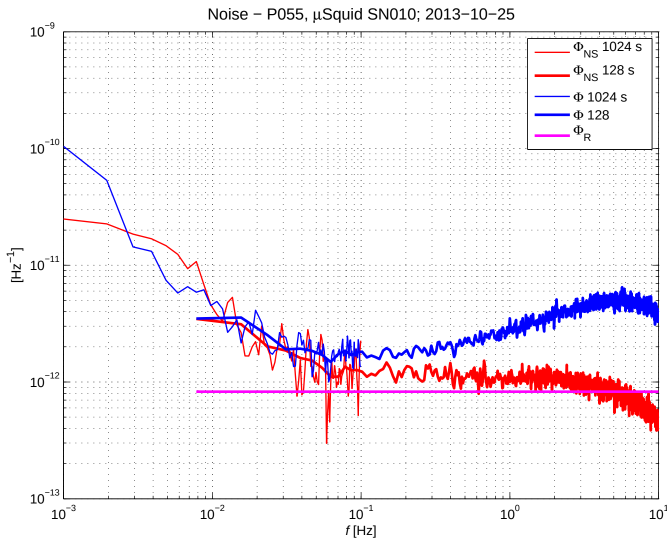

Figure 7: Concentration Noise Spectra. Same as Figure 4 but expressed in terms of a 100 % concentration.

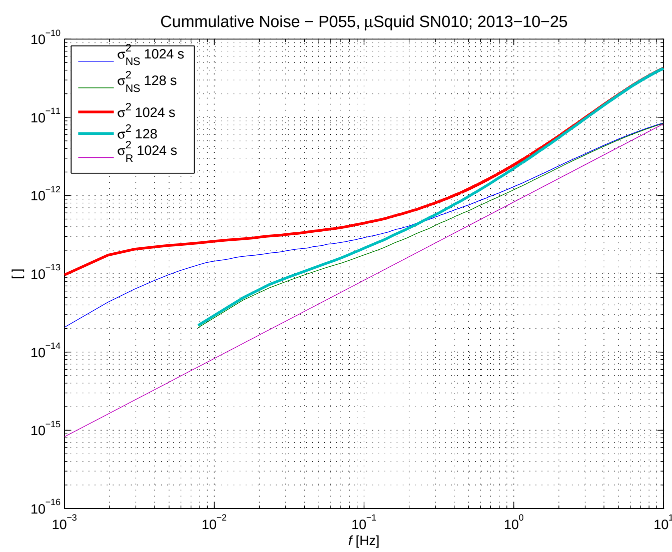

Figure 8: Cumulative Concentration Noise Spectra. Same as Figure 5 but expressed in terms of a 100 % concentration.

7 Discussion

7.1 Time- versus frequency-domain estimation

Many (but certainly not all) researchers using the eddy-covariance method of estimating the vertical flux of oxygen are uncomfortable with frequency-domain analysis. In fact, they usually calculate the covariance in the time domain by literally multiplying the de-trended measured vertical velocity with the de-trended oxygen concentration, point by point, and then compute the average value of this co-product over some interval of time (usually 1 to 10 minutes). If the flux is statistically steady, then the average value of the co-product is independent of time. The same result can be achieved in the frequency domain by integrating co-spectra. To demonstrate this equivalence, I have computed the noise signal variance using intervals of 1024 and 128 seconds for 1000 randomly chosen starting intervals of the noise time series with a sensor installed (green trace, Figure 3). The histograms of the standard deviation of the signals (Figures 9 and 10) agree with the cumulative variance (which is the square of the standard deviation) at 10 Hz (Figure 5). For example, the concentration variance (at 10 Hz in Figure 8) is 4×10^{-11} and so the standard deviation is the square root of this value, namely 6.3×10^{-6}. The value obtained by time-domain calculations is 5.3×10^{-6} with a 95 % confidence interval of 4.6 to 5.9×10^{-6} (Figures 9 and 10).

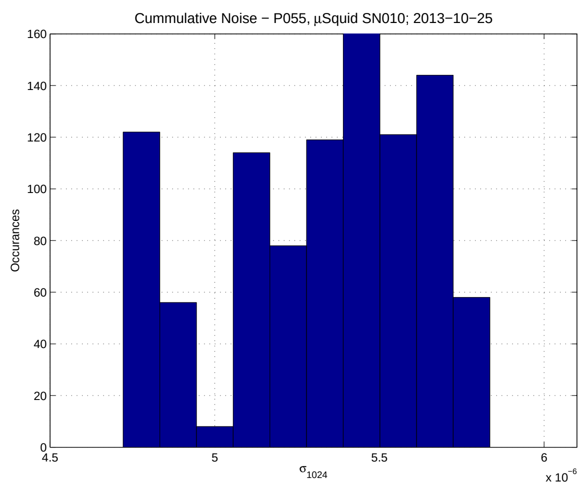

Figure 9: The distribution of 1000 standard deviation estimates using 1024 seconds of data with random starting points.

The advantage of frequency-domain estimation is the flexibility of choosing the band of interest so that one includes only the frequency range that has a meaningful environmental signal and excludes the frequency range dominated by noise. The result is an estimate that has a superior signal-to-noise ratio and the quantities estimated have less statistical uncertainty. For example, if the environmental signal is significant only over the frequency band of 0.01 to 1 Hz then the variance over this band is, according to the cyan trace in Figure 8, only 2×10^{-12} and the noise standard deviation is 1.4×10^{-6}.

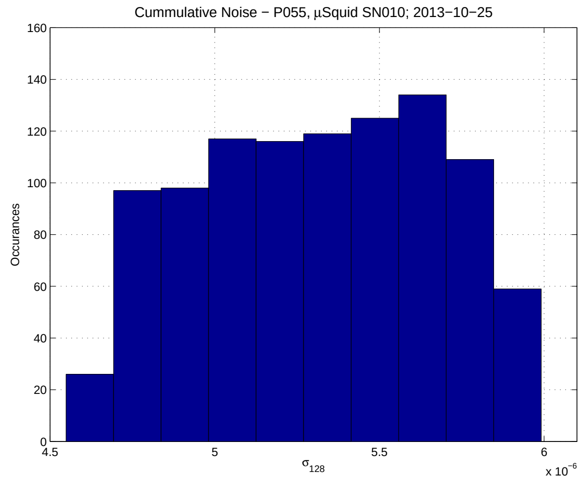

Figure 10: The distribution of 1000 standard deviation estimates using 128 seconds of data with random starting points.

7.2 Why is noise important to co-variance estimation?

It is frequently asserted that measurement noise is not important for the estimation of co-variance because the noise averages out to zero (Claudia Loraine, personal communication). This is true only in an asymptotic statistical sense. For example, if we represent the measured oxygen concentration by $\hat{C}(t) = C(t) + C_n(t)$, where C is the true environmental fluctuation of dissolved oxygen and C_n is the measurement noise, then the vertical flux of oxygen is

$$\overline{w\hat{C}} = \overline{wC} + \overline{wC_n} \tag{8}$$

where w is the fluctuation of vertical velocity and is assumed to be noise free, for simplicity. The vertical velocity (and any noise that it may contain) are statistically independent of the measurement noise, C_n, and, therefore, their co-product is statistically zero. So, statistically, $\overline{w\hat{C}} = \overline{wC}$. With progressively longer averages, the co-product $\overline{wC_n}$ gets progressively closer to zero, and vanishes asymptotically.

However, any two time series that are statistically independent and totally unrelated will still have a non-zero co-product if the number of samples used to form that co-product is finite[1]. This means that your estimates of the vertical flux of oxygen, $\overline{w\hat{C}}$ will fluctuate around the true value, $\overline{wC}$ because of the measurement noise. The fluctuations and, hence, your confidence interval is reduced by reducing measurement noise. Noise is reduced by using quieter electronics, a higher sampling rate, by selecting only the frequency range where the co-spectrum of w and C are significant, and by using a sensor that outputs more current per unit of oxygen signal.

End Of Document

[1] This is easily checked by generating two random and independent sequences and calculating their co-product using, for example, the Matlab function randn.