TIME CONSTANTS and DC GAIN

Remember, the output of the PSD contains many signals. Most of the output signals have frequencies which are either the sum or difference between an input signal frequency and the reference frequency. Only the component of the input signal whose frequency is exactly equal to the reference frequency will result in a DC output.

The low pass filter at the PSD output removes all of the unwanted AC signals, both the 2F (sum of the signal and the reference) and the noise components. This filter is what makes the lock-in such a narrow band detector.

Time Constants

Lock-in amplifiers have traditionally set the low pass filter bandwidth by setting the time constant. The time constant is simply $1/2\pi f$ where f is the -3 dB frequency of the filter. The low pass filters are simple 6 dB/oct roll off, RC type filters. A 1 second time constant referred to a filter whose -3 dB point occurred at 0.16 Hz and rolled off at 6 dB/oct beyond 0.16 Hz. Typically, there are two successive filters so that the overall filter can roll off at either 6 dB or 12 dB per octave. The time constant referred to the -3 dB point of each filter alone (not the combined filter).

The notion of time constant arises from the fact that the actual output is supposed to be a DC signal. In fact, when there is noise at the input, there is noise on the output. By increasing the time constant, the output becomes more steady and easier to measure reliably. The trade off comes when real changes in the input signal take many time constants to be reflected at the output. This is because a single RC filter requires about 5 time constants to settle to its final value. The time constant reflects how slowly the output responds, and thus the degree of output smoothing.

The time constant also determines the equivalent noise bandwidth (ENBW) for noise measurements. The ENBW is NOT the filter -3 dB pole, it is the effective bandwidth for Gaussian noise. More about this later.

Digital Filters vs Analog Filters

The SR830 improves on analog filters in many ways. First, analog lock-ins provide at most, two stages of filtering with a maximum roll off of 12 dB/oct. This limitation is usually due to space and expense. Each filter needs to have many different time constant settings. The different settings require different components and switches to select them, all of which is costly and space consuming.

The digital signal processor in the SR830 handles all of the low pass filtering. Each PSD can be followed by up to four filter stages for up to 24 dB/oct of roll off. Since the filters are digital, the SR830 is not limited to just two stages of filtering.

Why is the increased roll off desirable? Consider an example where the reference is at 1 kHz and a large noise signal is at 1.05 kHz. The PSD noise outputs are at 50 Hz (difference) and 2.05 kHz (sum). Clearly the 50 Hz component is the more difficult to low pass filter. If the noise signal is 80 dB above the full scale signal and we would like to measure the signal to 1% (-40 dB), then the 50 Hz component needs to be reduced by 120 dB. To do this in two stages would require a time constant of at least 3 seconds. To accomplish the same attenuation in four stages only requires 100 ms of time constant. In the second case, the output will respond 30 times faster and the experiment will take less time.

Synchronous Filters

Another advantage of digital filtering is the ability to do synchronous filtering. Even if the input signal has no noise, the PSD output always contains a component at 2F (sum frequency of signal and reference) whose amplitude equals or exceeds the desired DC output depending upon the phase. At low frequencies, the time constant required to attenuate the 2F component can be quite long. For example, at 1 Hz, the 2F output is at 2 Hz and to attenuate the 2 Hz by 60 dB in two stages requires a time constant of 3 seconds.

A synchronous filter, on the other hand, operates totally differently. The PSD output is averaged over a complete cycle of the reference frequency. The result is that all components at multiples of the reference (2F included) are notched out completely. In the case of a clean signal, almost no additional filtering would be required. This is