


Lock-in Amplifier [Notes]

What does it do?

The lock-in amplifier is a device used to extract very weak signals from noisy measurements. By knowing the exact frequency of the signal, it can zero in on that specific part of the spectrum, effectively isolating and amplifying the signal despite the surrounding noise.

In particular it's more proper to think of the lock-in as a spectrum analyzer. This device stimulates a physical system (or Device Under Test DUT) with a periodic signal, the system responds with a periodic signal of its own + noise. The job of the lock-in is then to compute the Fourier Transform of this waveform and extract the first harmonic amplitude and phase.




How does it do it?


I've chosen to write 2 explanations here, both are in my opinion necessary to fully understand a lock-in however one is more suitable to understand the lock-in from a signal processing point of view, the other is a more proper explanation of it's theory of operation but it's a bit more in depth.

The lock-in from a signal processing standpoint or The short explanation


A Lock-in amplifier can be abstracted to just a couple of blocks:



The stimulation signal called “reference” ($V_r(t)$) comes out of the lockin and stimulates a physical system, the physical system responds with an “input signal” ($V_s(t)$) that is fed into the lock-in. In the diagram above we can see that the circuit is duplicated save for a 90° phase shift, for now we'll only evaluate the top half. There we can see that the signal is “mixed” with a reference signal. This process will be explained more in depth in the following section but for now we'll believe that the mixer takes 2 signals and outputs a third. Below we can find the frequency space of the signals before and after the mixer:




It takes signals f_{sig} and f_{REF} and outputs a signal with 2 components: $f_{sig} - f_{REF}$ and $f_{sig} + f_{REF}$. However this case assumes that the input signal coming from the experiment is purely sinusoidal (i.e. defined solely by a peak at f_{sig}), this is not the real case, the frequency spectra of our signal is in fact plagued with noise at all frequencies. However the Fourier Transform tells us that this spectrum of noise + signal can be decomposed into many different peaks, each at a different f_{sig} . Knowing this we can see that the noise spectrum will be shifted by the mixer but our signal will be shifted in a special way: In the case where the signal is a direct response of the stimulation at f_{REF} it's spectrum component after the mixer will be shifted to 0Hz and $2 f_{REF}$. That is, mixing our noise + signal spectrum with f_{REF} allows us to separate the signal from the noise by shifting out signal peak to DC and everything else into AC. Knowing this we may separate both by applying a low pass filter to the mixer output which will preserve only the DC component, this component is the output of the lock-in and it is proportional to the physical signal amplitude.



However the signal experiences a delay when traveling through our system, this delay can be also understood as a phase when dealing with frequencies and some phases may collapse the output of the lockin to 0, for reasons that will be explained in the next section, we can combat this by computing this process twice with a duplicated circuit but with a signal delayed by 90° .

We may do one final leap of faith to better understand the role of this instrument in our signal chain.




Initially we described the mixers from the point of view of the lock-ins output, that is the whole spectrum is shifted by $-f_{REF}$ bringing our signal at $f_{sig} = f_{REF}$ into DC, however it's perhaps more useful to think from the reference frame of the input and interpret the mixers as frequency shifting our Low pass filter from DC into f_{REF} (see image d above) Our low pass is now no longer allowing to pass only spectra between 0Hz and $f_{cutoff} = \frac{1}{2\pi\tau}$, it is now a band pass filter centered in f_{REF} and allowing to pass frequencies at either side by $\pm f_{cutoff}$. This reduces the whole complicated signal chain inside the lockin into one single understandable block, a really fine band pass filter.

Note: While it's fair to understand the lockin as a bandpass at the input signal we still have to take into account the frequency shift happening at the output. For example let's take a lock-in configured to $f_{REF} = 1\text{KHz}$, low pass cutoff frequency set to $f_{cutoff} = 1\text{Hz}$ (or time constant $\tau = \frac{1}{2\pi \cdot 1\text{Hz}} = 159\text{ ms}$) this means from the input point of view we can expect to take only the frequency components from 999 to 1001 Hz but we should not expect an output from 999 to 1001 Hz like we would in a bandpass, we should expect to see an output from DC to 1 Hz, the amplitude characteristics of the spectra however will indeed be the same.

The lock-ins theory of operation or The longer explanation:

The following is the block diagram for the SR830 lock in amplifier, it may seem daunting at first but through the following explanation we will come to see how the signal processing that's taking place is exactly a Fourier Transform computation.



Lock in amplifiers in essence compute the Fourier Transform for a certain harmonic of f_{REF} . To demonstrate how let's look at the formula for the FT:

$$\hat{f} = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ft} dt$$

If we expand the complex exponential into cartesian coordinates the picture is made more clear

$$\hat{f} = \int_{-\infty}^{\infty} f(t) [\cos(2\pi ft) - i \sin(2\pi ft)] dt$$

We may interpret the FT as a measure of similarity, the FT amplitude plot represents how "similar" our function $f(t)$ is to a certain function, as a first approximation we'll compute the similarity to functions of type: $\cos(2\pi ft)$ and later $-i \sin(2\pi ft)$.

The FT formula above is taking a couple of computation steps that a lock-in will try to take too, first it's evaluating how similar our signal $f(t)$ is to a cosine function of frequency f by multiplying both of them, to understand how this is a measure of similarity let's start with an example. take our signal to be another cosine of an arbitrary frequency:

$$f(t) = V_{sig} \cos(2\pi f_{sig} t + \theta)$$

And instead of computing the full spectrum of the FT let's evaluate only the similarity to a cosine function with a particular frequency f_{REF} :

$$V_{REF}(t) = \cos(2\pi f_{REF} t)$$

For which after multiplication we'll get:


$$f(t)\cos(2\pi f_{REF} t) = \frac{1}{2}V_{sig} [\cos(2\pi t[f_{sig} + f_{REF}] + \theta) - \cos(2\pi t[f_{sig} - f_{REF}] + \theta)]$$

This step is called the “demodulation” and its represented by a multiplication sign or a “mixer” in the block diagram.

As explained on the previous section after multiplying 2 signals we get an output signal with 2 components in the frequency space, one at $f = f_{sig} + f_{REF}$ and another at $f = f_{sig} - f_{REF}$. It is now clear that in the special case where $f_{sig} = f_{REF}$ we get a signal centered in 0Hz plus another at $2f_{sig}$.

This is the main principle on how both Fourier Transforms and Lock-in Amplifiers “evaluate similarity” of an input signal to a given frequency: After demodulation the only signal with a DC component will be that where $f_{sig} = f_{REF}$ and thus all we have left to do is eliminate all signals that do not have a DC component, to do so we take the integral of the output since any periodic function that does not have a DC component will average to 0 and it's integral from $-\infty$ to ∞ will be 0 as well. This is done though low pass filtering the output of our mixer, it is perhaps not a well known property of low pass filters but any frequency that is well above their cutoff frequency will be integrated at the output.

(Minor pet peeve: This step is sometimes described in the literature as the lockin taking the average of the mixer output but that assertion is mostly incorrect since an average does not have the same frequency response as an integration/low-pass, it's also inaccurate when comparing with the formula of a Fourier Transform)



- a) Input mixer: Special case where $f_{sig} = f_{REF}$
b) Output mixer: Special case shows DC and AC component
c) Input mixer: Every other case, none of those show a DC component
d) Output mixer: Every other case will show only an AC component

We have now explained how this mathematical/hardware tool is able to evaluate if a signal is the same frequency as another known frequency f_{REF} , we should now clarify that the signals we plan on measuring carry all sorts of components but the Fourier Transform guarantees that only those that correspond to the reference signal will survive.

Of course we have yet to answer why perform this computation twice for a cosine and a sine. This is because our signal may have a phase that is different to the phase of our cosine: $f(t) = V_{sig} \cos(2\pi f_{sig} t + \theta)$, and there are phases for which the multiplication of both cosines is 0 or anywhere in between, in fact if you do the math the surviving term after the integral is actually:

$$\hat{f} = \frac{1}{2} V_{sig} V_{ref} \cos(\theta)$$

To solve this dependency on θ we may make use of the fact that the sine is “orthogonal” to the cosine, that is it is impossible for a certain phase delay to simultaneously collapse the multiplication and integral to 0 for both the sine and the cosine. This is the reason why we see 2 parallel branches in the block diagram, one computes the cosine part of the formula and the other delays the reference by $\theta = \frac{\pi}{2}$ to compute the sine part of the equation.

Therefore we end up with 2 outputs, X and Y in the block diagram, however we may now come back to the very first step we took after laying the formula for the FT where we expanded the complex exponential into 2 cartesian terms and realize we can make use of the complex polar coordinates and transform our outputs into a more useful representation where:

$$R = \sqrt{X^2 + Y^2}$$

$$\theta = \text{atan2}(Y/X)$$

Now R properly represents the correct value of our first harmonic in the Fourier Transform without dependencies on the phase delay and θ represents the delay between the reference signal and the received signal.

One final clarification is the formula for the FT integrates time from $-\infty$ to ∞ but this is obviously impossible, we should expect this because the outputs of a proper FT are delta functions of infinite height centered at the frequencies present in our spectra. Something that we'll never see when integrating for finite periods of time. Instead we can modify the FT formula to account for the finite integration time:

$$V_{out}(t) = \frac{1}{T} \int_{t-T}^T f(s) e^{-i2\pi f_{REF}s} ds$$

This integration time is set by the real life characteristic of our low pass filter, our low pass filter has a time constant (settling time) that is inversely related to frequency:

$$\tau = \frac{1}{2\pi f_{cutoff}}$$


Thus we arrive at our first practical constraint, if we want to reject more noise we have to integrate for longer periods of time (i.e. lower our cutoff and reject more noise), this should make sense since it approximates our measurement to the ideal definition of a FT.

How to use one

TO DO

How to set up the low pass filter?


The following is the frequency, phase and step response of a cascade of n filters (*commonly the amount of filter stages is specified in the lock-in panel not by n but by the slope of it's attenuation: $n = 1 \rightarrow 6\text{dB/oct}$, $n = 2 \rightarrow 12\text{dB/oct}$, $n = 3 \rightarrow 18\text{dB/oct}$, $n = 4 \rightarrow 24\text{dB/oct}$*), all of them share the same time constant τ which is the case for a lock-in amplifier, a couple things to note is how the settling time at the step response delays, how the cutoff frequency (*strangely imo, will have to check*) changes too. (Note: The phase should be related to n however I'm not seeing that, I guess the lock-in corrects for it).



How to setup the SRS860's Gain:

Careful, in the SR860 the actual Gain that may affect the actual signal to noise ratio for the measurement is controlled by “*Input Range*”, “*Sensitivity*” is more akin to an artificial zoom into the digitized waveform that does nothing to the measured value.

See the following block diagram for the SR860, looks complicated but just focus on the gray area, those are the processes that are happening in the digital realm, that's where you'll find the Sensitivity control, now, everything outside that is analog and at the very beginning of the signal capture you'll see the input range control that set's up the analog Gain.



To put it simply and into the perspective of an optics scientist, imagine the lockin as a digital camera, the focusing lens would be an analogy for the Input Range, allowing you to capture a crisper image, while the Sensitivity would be more akin to zooming into the digital image after taking the picture. The focus will allow you to capture more relevant information while the digital zoom adds nothing of value to the signal.

My thinking is the sensitivity set up is there to allow the user to capture better data directly from the output BNCs (Ch 1, Ch 2...), if you are gonna place an oscilloscope there it does pay to artificially raise that signal. But we don't care about that since we are capturing the digitized signal through RS-232.

How to connect the SRS860 to a computer through Ethernet:

A ChatGPT guide:

1. Connect the Devices

Use an Ethernet cable to connect your PC directly to the lock-in amplifier.
Ensure both devices are powered on and ready.

2. Configure the Lock-In Amplifier

On the lock-in amplifier's settings screen by pressing and holding Calc System:
IP Address: Set to 192.168.0.4.
Subnet Mask: Set to 255.255.255.0.
Default Gateway: Leave blank or set to 0.0.0.0.
Save and apply the changes.

3. Configure Your PC's Static IP Address

Open Network Connections by pressing Win + R, typing ncpa.cpl, and pressing Enter.
Right-click your Ethernet Adapter and select Properties.
In the list, select Internet Protocol Version 4 (TCP/IPv4) and click Properties.

Set the following:

IP Address: 192.168.0.5 (ensure the last digit differs from the lock-in amplifier's IP).
Subnet Mask: 255.255.255.0.
Default Gateway: Leave blank or set to 0.0.0.0.
Preferred DNS Server: Leave blank or use 8.8.8.8 (optional).
Click OK and close all windows.

4. Verify the Connection

Open Command Prompt by pressing Win + R and typing cmd, press Enter.
Test the connection by pinging the lock-in amplifier typing:
ping 192.168.0.4
If you see replies, the connection is working.

5. Test with Python notebook

Questions:

1) What does the output of a lock-in mean when I chop a DC signal?

Let's say I point a continuous source of light like a laser into a photodiode, this photodiode should generate a DC current but it may be too weak to measure so I employ a lockin amplifier and a chopper, the chopper, in essence, modulates this DC source into a square wave by alternatively blocking the light path at a regular interval, the frequency from the chopper is fed into the lockin as a reference. This is a pretty standard way of measuring optical quantities.

However there is something about this measurement that irks me and it goes as follows: If the light is chopped into a square wave then the DC current I wanted to measure initially will be chopped into a square wave as well, distributing the "energy" of the signal among all of the many harmonics a square is made of, if I then use the lockin to measure the amplitude of the first harmonic the amplitude of this harmonic is cannot in my mind be the same as the DC current level I was trying to measure in the first place, right? It might be proportional but not an absolute measurement right?

Answer?: I believe what should be done here is expanding the square wave into a Fourier Series (being careful to setting an offset such that the light power doesn't go negative ofc) then finding out the relationship between the peak value of the square wave (which was the DC value before chopping) and the first coefficient in the series which I think correlates to the amplitude of it's peak. (not really sure you can use fourier transforms and series interchangeably like that tho)

However we won't be measuring DC values which leads me to the next question:

- 2) **What kind of waveform are we trying to measure?** (taking chopping pump probe delays and TGFWM into account).

Answer?: Let's say that Cris wants to capture the evolution of a signal over time, for example the decay of signal power as the sample's material relaxes. This evolution will be a modulation of amplitude $m(t)$ in the reference signal:

$$s(t) = m(t) V_s \cos(w_{REF} t + \phi)$$

If we understand a lock in system as a bandpass centered around w_{REF} with bandwidth equal to $2 f_{cutoff}$ then we'll see that the temporal evolution of the signals amplitude will be captured only if it's within the vicinity of $f_{REF} \pm f_{cutoff}$, that is if the amplitude modulation signal $m(t)$ "survives" the low pass filtering. This question then doesn't make sense for a pump probe experiment since the pump probe method exists solely to study super fast phenomena, unless we use a super high f_{REF} in the order of the phenomena we try to examine.

Which leads me to believe that the signal we are going to capture is the super fast signal we want to study integrated from a certain time step, set by the delay line, up until it's extinguished and I doubt that this period is longer than the chopper time step so I'm guessing we'll have many pulses of laser in between each "chop" and then no pulses when blocking the laser. This all feels like a hard problem to extract an absolute measurement but it probably works as a relative measurement.

I think I understand why we chop one out of 2 pulses. The way I understand it chopping will give us a probe signal after pumping followed by an unpumped probe. This gives a pulse train that has some envelope at the chopping frequency (half of the laser trigger). The higher the effect of pumping the bigger the difference between pump and unpumped and thus the bigger the envelope and the bigger the sub-harmonic. If on the contrary pumping served no purpose we would all pulses would look the same and there would be no envelope so we would find nothing but noise when evaluating the first subharmonic frequency.

References

- All of the papers and books in the parent folder at this Google Drive
- TSP SR830 [[video](#)]
- TSP lock-in demonstration [[video](#)]
- 3B1B Fourier Transform [[video](#)]

TO DO: Watch these:

https://www.youtube.com/watch?v=Yqs-4TKGjjQ&t=2s&ab_channel=DallinDurfee

<https://www.youtube.com/watch?v=kNnc-IDHtrM&t=19s>

<https://www.youtube.com/@analogwings/videos>

https://www.youtube.com/watch?v=h0_qrLpeT3k&ab_channel=snarayan007