

The measurement of loop gain in feedback seismometers

Brett M. Nordgren

April 9, 1999

Rev. April 5, 2010

brett3nt@bnordgren.org

Introduction

In reading the messages coming through the PSN-L mailing list, I have noticed that on several occasions people were having problems with misbehaving feedback loops. On occasion, it appeared that promising designs or design features were being abandoned or seriously modified just to get rid of oscillations when in fact that may not have been necessary. It is common for oscillations to occur in new designs or in old designs after being modified. The experienced feedback designer simply expects that and plans the evaluation and testing accordingly.

The difficulty with cut and try design techniques on feedback loops, is that to some degree the various components of the design all interact. If changing one parameter results in oscillation, it still cannot be said that that parameter *caused* the oscillation. It is quite possible that some modest changes elsewhere in the loop may solve the problem perfectly well. Trying to get a loop to work, blind, is one of the more frustrating activities I can think of.

All that having been said, what do we need to do to understand and correct loop stability problems? Here we are talking about *loop gain*, the amplification around the feedback loop, which measures the "strength" of the feedback as a function of frequency (See: feedback.pdf, loop3.pdf and bbab-028.pdf at <http://bnordgren.org/seismo/>), as that is what determines the effectiveness of the feedback as well as any tendency for the circuit to oscillate. Regarding oscillation, the interesting loop characteristics are at the relatively higher frequencies, perhaps a decade below and above the gain-crossover (loop gain = 1) frequency, which is 37 Hz in one design. For that design, it means that the loop characteristics at frequencies from 3 to 400 Hz are the most important, with 400-2000Hz being of interest if mechanical resonances are suspected. Although low-frequency parameter changes may affect the loop, if they are creating problems it is likely only because they are in turn affecting the loop at the higher frequencies, usually just by raising or lowering the high frequency portion of the loop gain curve as a whole.

Measuring loop response

Measurements in the 1 – 2000 Hz range are generally quite easy to make, and any feedback seismometer design project should include such tests. In the books, the way loop gain is measured is to break the loop at *any* point, inject a sinusoidal signal into the downstream free end and measure the returning signal coming back from the upstream free end. Their ratio is the loop gain and their phase difference is the loop phase. The signals injected and measured can both be currents, voltages, forces, accelerations--whatever is appropriate to the place where the loop was broken. A secondary consideration is that the broken ends should be terminated with passive impedances to match what they were seeing when the loop was connected.

In practice, this measurement technique is very difficult to use successfully.

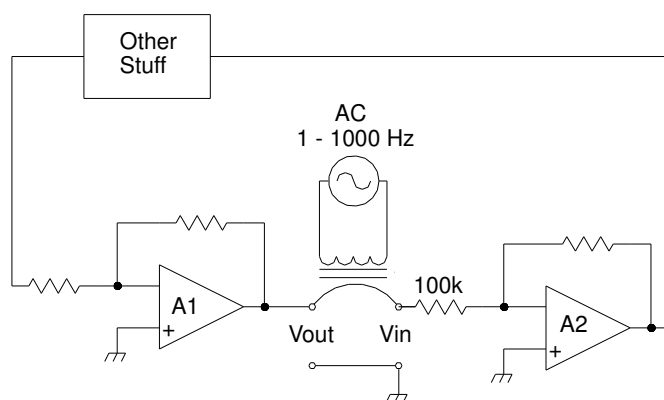
Most feedback loops are employed to help maintain a fixed operating point for the system. In the STM-8, the position feedback is acting to hold the boom centered at the desired level. The critical characteristic for that to work is that there is some reasonable amount of loop gain at DC. In such a loop, breaking it may result in the instrument losing the ability to maintain its balance, making measurements difficult or impossible. What is needed is a technique for measuring loop gain while the feedback loop remains closed and operating normally. Such a technique was analyzed in the Jan-Feb 1963 issue of the *Hewlett-Packard Journal*: “A quick convenient method of measuring loop gain” by Philip Spohn. see:

<http://www.hpl.hp.com/hpjournal/pdfs/IssuePDFs/1963-01.pdf> This method made use of their [then] new 310A Wave Analyzer, but it will work with any system in which an appropriate sine wave source, and means for measuring small AC signals are available.

The Inserted Voltage technique

Spohn observed that an operating feedback loop may contain several natural break points where a downstream element presents no significant load to an upstream element. For voltages, this would be a place where a low impedance device, say an op-amp output, is driving something that has much higher impedance at the frequencies of interest. Or for currents, it would be a place where the input impedance of a downstream element is very low compared with the very high output impedance of the upstream element. It is important to observe here, that the measured loop gain function will be identical no matter where in the loop it is measured, and regardless of what type of signals are being measured. So we have considerable freedom to select the point in the loop at which we will choose to measure the loop gain.

In general, it is easiest to work at a point where the signal levels are highest, where the loop will be least sensitive to noise signals which will unavoidably get injected during the measurement. Also we need to use a point where the entire loop signal is passing through one path. In the P/D/I configuration, it will not work to make the measurement on only one of the three parallel branches. Spohn's technique consisted of inserting a voltage signal into the active loop and then measuring its relative size, immediately upstream and downstream of the insertion point. The downstream signal may be considered to be the input to the loop V_{in} while the signal measured on the upstream side is the loop output V_{out} . The amplitude ratio and phase difference of those signals accurately reflect the loop gain and phase at the measurement frequency.




Let us look at making the measurement with an inserted voltage. Here the loop break is assumed to be at an op-amp output (low impedance) driving a 100k resistor (high impedance). First, we will need a floating signal source. Spohn's idea was to take a passive AC clip-on current probe and use it in reverse, applying a voltage to its multi-turn secondary and clipping it onto a (1-turn) wire link in the feedback loop. This inserts a small AC signal between the

points labeled V_{out} and V_{in} . The loop gain is then calculated from measurements made of V_{out} and V_{in} to ground. At lower frequencies, the signal amplitude at which the probe core saturates

is quite low, and it may be necessary to use several turns linking the probe to get the inserted voltage high enough for a good measurement, and even that may not be enough.

While the AC probe technique can work well at kilohertz frequencies, at the frequencies of interest in seismometers it might not. However, an excellent method of voltage insertion which will work down to DC is to use an instrumentation amplifier set up for unity gain and wired into the loop to insert the test signal. In this method we are making use of the common-mode rejection capability of the in-amp to create a floating AC voltage source. In making an experimental feedback seismometer, I would probably build in one of these as a part of the circuit, and consider it time and money well spent.



If we are able to insert reasonably high-level signals, an oscilloscope is the ideal measuring tool, preferably one having two traces and good sensitivity down to DC. A blind AC voltmeter will work at higher frequencies, but it is still imperative to know that the inserted signal is not saturating the feedback loop, rendering the measurements meaningless, and also to know that you're not just measuring hum or noise voltages. With a scope viewing the two signals, overloading is easy to see, allowing the drive voltage level to be reduced until both sine waves are clean. At the lower frequencies, where the loop gain is still moderately high, this is a very sensitive way of detecting loop overloading, as any slight distortion will be evidenced by a gross distortion in the V_{in} waveform. The scope will also let the phase shift be observed, which is useful information.

At the lower frequencies, assuming that you can find a very low frequency signal source, a two channel data acquisition system, perhaps using the same setup planned for recording seismic data, might be used. This will give better measurements at lower frequencies than an oscilloscope, but since you are working in the blind you won't know until the data can be plotted whether the drive signal level was appropriate for getting a good measurement while still not overloading the loop.

Measuring loop gain

To make the measurement, connect up the scope, one channel to V_{in} and the other to V_{out} . Set both to the same fairly sensitive range, DC coupled if possible. Set the signal frequency to the expected gain-crossover frequency. Slowly raise the signal generator level and/or change the scope sensitivity until both channels are showing good sine waves, without distortion. You will find that making these measurements involves repeatedly re-adjusting the drive signal level and the scope input ranges to get clean, observable waveforms. Now tune the frequency until V_{out}

and V_{in} are identical in amplitude. That frequency is the gain-crossover frequency, and the phase difference +180 degrees is the *phase margin* of the loop which should also be recorded.

It can be confusing to properly interpret the phase and it may be helpful to invert the display polarity of the V_{out} signal, but it is almost certain that V_{out} will be lagging V_{in} by something between 90 and 180 degrees. Anything much beyond -140 degrees at gain crossover is getting a bit close to oscillating and is probably grounds for making some adjustments to the circuit. If that is the case, you'll probably also notice that the scope traces sharply increase in size as you tune through the gain-crossover frequency and you may have to turn down the signal level to get them back in range. Significant peaking at gain-crossover is a symptom of an almost-oscillating feedback loop.

The loop phase can also be easily computed from measurements of the three signal amplitudes, V_{gen} , V_{out} and V_{in} , where V_{gen} is the amplitude of the inserted voltage. The loop phase $= \cos^{-1} [(V_{out}^2 + V_{in}^2 - V_{gen}^2) / (2 V_{out} V_{in})] - 180 \text{ deg.}$, which will be generally be a negative angle. When examining the loop at the gain crossover frequency, the frequency where $V_{out} = V_{in} \equiv V$, the loop phase will be $\cos^{-1}(1 - (V_{gen}^2 / 2V^2)) - 180$ and the phase margin will be just $\cos^{-1}(1 - (V_{gen}^2 / 2V^2))$. A phase margin which is greater than 40 degrees or so, should be adequate to avoid serious gain peaking at gain crossover and the related risk of loop oscillation.

Make from four to ten measurements of V_{out}/V_{in} at frequencies down to 1/10 of gain-crossover and a similar batch in the frequency decade above it. If you find a frequency at which the amplitudes suddenly peak as you tune through, you're probably observing a mechanical resonance. Resonances can be a serious problem at frequencies above gain crossover. If they peak too high they can take a loop gain which is below unity and falling nicely with frequency, and put it back up through unity, creating a second gain crossover frequency, one which will almost certainly oscillate. This is also a quick test to observe how well your attempts at damping resonances are working. Since the STM-8 uses a 600 Hz sensor drive frequency, I would expect that any measurements above 300 Hz would not be meaningful, though it might be interesting to see what happens.

The most useful way to display the measured data is to plot, on paper or on the computer a log-log graph of loop gain (V_{out}/V_{in}) vs frequency. Aside from possible resonances, the plot should be a smoothly falling curve. If it is not smooth, the likely problem is that the loop was overloaded during one or more of the measurements. Now note the slope as the plot descends through gain crossover. In the vicinity of gain-crossover the slope of the curve should not be a lot steeper than -1 (i.e. 10:1 per frequency decade). -1.5 is probably OK but if it is approaching a -2 slope (100:1 per decade), the loop will be very close to oscillating. This is because in most circuits, a -2 slope is associated with a phase shift of -180 degrees. Occurring at gain-crossover, that is the criterion for oscillation. If the slope is -1, the phase shift will be roughly -90 degrees, which is extremely comfortable. Likely you will measure something in between.

Dealing with problems


The cure for a marginal loop is to examine those circuit components affecting the loop in the frequency region around gain crossover (within 1 decade in frequency above and below) and modify them to make the slope lower in that region. That is often not too hard to do, though

each circuit will need its own approach. In general, the cure will best be implemented in the higher frequency parts of the circuit, near gain crossover, rather than in the main loop elements. The process of improving loop feedback stability by modifying transfer functions is termed “loop equalization”, and the components added for this purpose are called an “equalizer”. There is much more which could be said on this. Equalizing loops is how feedback designers earn their pay.

You will notice that I haven’t focused here on measuring and plotting loop phase shift. That is partly from having observed that the slope and shape of the loop gain curve quite reliably implies the phase shift. The only case in which it might not would be when encountering lattice or bridge type circuits, so-called non minimum phase networks, which are only rarely seen in feedback loops. Since design modifications to improve loop stability are best visualized and implemented as modifications to transfer functions, one gets into the habit of looking mostly at gain and not phase. However, it is an excellent idea to quickly check the phase shift at the gain-crossover frequency (the phase margin) to verify that it is roughly as predicted, or to confirm that loop equalization has been successful or that changes made for other reasons haven’t seriously reduced loop stability.

The Injected Current technique

We have looked at voltage insertion techniques to measure loop gain. Injecting a current and measuring the upstream and downstream currents might also be possible. The current injection source can be an AC signal generator connected to the circuit through a high value resistor. Usually the maximum current that can be injected without over-driving the feedback loop will be quite small. If we are measuring in the decade above and below gain crossover, some measurements will likely have to be made of currents which are more than 10x smaller than the injected current. Since current probes do not have the required sensitivity, the measurements must be of the voltage drop across some series resistor. If a resistor must be added for that purpose, it has to be small enough to not affect the circuit, implying even smaller voltage measurements. In addition, those measurements must usually be made with differential inputs, as they are often floating with respect to ground.



Nonetheless, on rare occasions a point may be found which allows this kind of measurement to be tried. The circuit we looked at before is a case in point. If we choose to break the loop between the 100k resistor and the summing point of A2 we now see the required high impedance upstream and low impedance downstream.

We note that at all frequencies of interest, the current into the summing point of A2 is mirrored (negatively) by the voltage across

its 10k feedback resistor. Since the summing point is at virtual ground, the measurement can be made with respect to ground, and $I_{in} = -V_2 / 10k$. However, across the 100k resistor the voltage must be measured differentially. Checking the scope probe input capacitances and resistances, their impedance needs to be high relative to the A2 summing point and to the A1 output, which they certainly should be at any frequencies of concern to us. And the scope's common mode rejection must be adequate. Here, all the conditions for a successful current-based measurement seem to be fulfilled, with $I_{out} = V_1/100k$. We inject the current through a 1 Meg resistor into the selected break point and observe that loop gain $= I_{out} \div I_{in} = -V_1/V_2 \times 0.1$

Some Final thoughts

I should briefly touch here on how to handle a loop that is actually oscillating. Often the gain in some part of the loop can be temporarily reduced (at all frequencies) enough to get the loop to settle down. Then the measurements can be made, bearing in mind that the actual loop gain numbers are higher than the measurements by the factor by which the gain was reduced. Now, using the information obtained, an equalization network can be formulated which, if successful, will eliminate the oscillation once the gain is returned to normal.

Since most of my feedback experience is with systems which are not sensitive to vibration or tilt, I'm not sure how much the presence of the operator will add to the difficulty of the measurements, although it is true that you can be working here with relatively high-level signals. I can only suggest that if a really big earthquake comes along while you're measuring, you're probably going to have to wait awhile. [Author's note: Since this was written, that has happened on one occasion, when an earthquake arrived during a measurement run, momentarily causing considerable confusion.]