Delivering Value Through Innovation

DSP Measurement of Under-Settled Signals

by Eric Rosenfeld, former Staff Scientist for LTX Corporation (now LTX-Credence)

Abstract

This paper explores the effects of settling transients on DSP sinewave measurements and techniques that allow measurements to begin before settling is complete. It shows the effects of many settling transients can be virtually eliminated.

Introduction

In many sine-wave test situations, test techniques must be optimized to perform tests in the minimum amount of time. Specialized DSP techniques such as coherent measurement are frequently used as an alternative to windowing non-coherent measurements.

One of the major test time components for low frequency band-pass devices (such as telecom SLICs) is the time required for the output transient of the device-under-test (DUT) to settle prior to measurement. In many cases, if required measurements can be completed before the device has completely settled, significant amounts of test time can be saved. This paper explores the effects of settling transients on sinewave measurements, and techniques for limiting such effects. One such technique, spectral interpolation, has been proposed as a means of accomplishing this. Standard DSP window techniques can also be used to limit the effect of transients on unsettled measurements.

This paper explores the relative merits of spectral interpolation and windowing. It also presents some general analysis on using windows in the context of coherent measurements of undersettled signals.

A DSP Measurement System

A typical DSP measurement system is shown in Figure 1. The system consists of a stimulus source that generates signal samples at a controlled rate. Repetitive signals are generated by repeating the same sample set as many times as is required. The measurement subsystem collects samples of the device output. Usually, the sample set collected with the measurement subsystem is analyzed using Fourier analysis techniques.

The specific details of the analysis will depend on the relationship between the sample set size and sample frequency of the source and measurement sub-systems. If the total time spanned by ‘n’ repetitions of the source sample set exactly matches the time spanned by the measurement collection interval (the product of the number of measurement samples multiplied by the time between measurement samples), the overall measurement system is coherent.

In order for the DSP source to produce pure sinewave components, the sample set used for the DSP source must be chosen so that it contains a whole number of cycles. If this constraint is not satisfied, a signal discontinuity will occur each time the signal pattern repeats, resulting in unintended signal components.

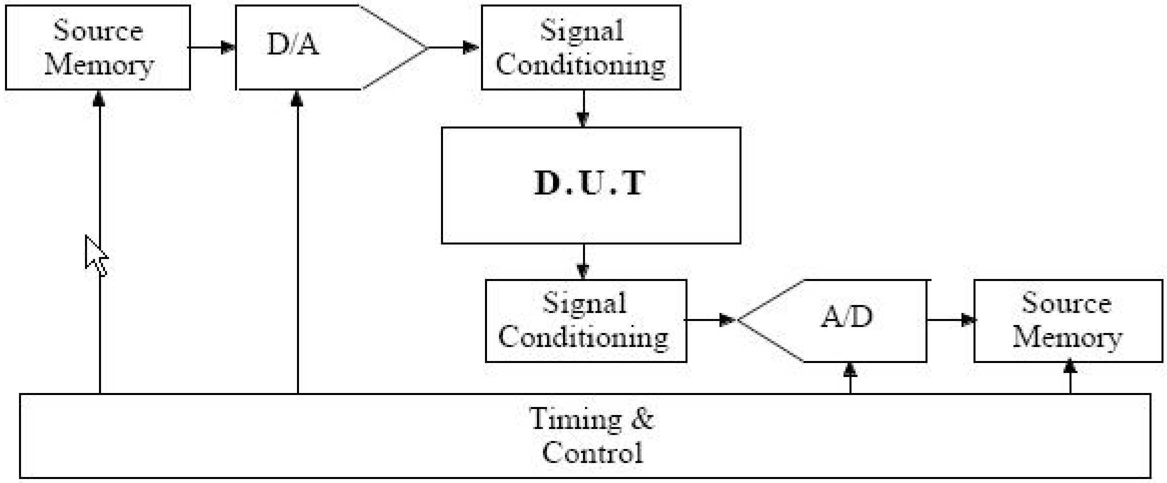

Figure 1 - DSP Measurement System

Non-Coherent Measurement

Standard DSP analysis techniques such as the Direct Fourier Transform, or the Fast Fourier Transform assume that the collected data is periodic over the measurement interval. In particular, these techniques analyze the signal as a sum of sinewaves at harmonics of the measurement period. If the measured signal does not exactly repeat over the measurement interval, Fourier analysis of the collected measurement data will yield signal components at frequencies far from the actual frequency of the measured signal.

The technique of windowing is frequently used in such situations to minimize the effect of the discontinuity between the last and first sample in a non-coherent sample set.

Coherent Measurement

A coherent measurement is setup so that the measurement interval exactly spans a whole number of cycles of all the signal components. When the measurement period satisfies this constrain, signal components fall exactly at analysis frequencies. There is no ‘leakage’ of any signal component into other analysis frequencies.

Because of the exact relationship between signal components and analysis frequencies, windowing is neither necessary nor used with coherent measurements.

Non-Coherent Interference

Even though the stimulus and measurement for a DSP test may be carefully controlled so that the measurement period spans a whole number of signal cycles, interfering signals are frequently not synchronized. Two common sources of interference, the power line (mains) frequency and settling transients are never repetitive with the measurement period.

Because interference is frequently not synchronized, it can have a significant effect on a coherent sine wave measurement. Settling transients are not periodic. As a result, the settling transient will appear as a damped exponential superimposed on the desired test stimulus. When Fourier analysis techniques are applied to a measured signal containing such a transient, frequency components across the entire analysis bandwidth will appear. To a first order, the interference introduced at most analysis frequencies will be proportional to the magnitude of the difference between the first and last measured sample due to the transient, and inversely proportional to the analysis frequency. The details of the transient, as long as the transient varies slowly over the measurement interval, will not strongly affect the spectrum of the interfering components.

There are at least two different techniques that can be used to for minimizing the effect of such interference. These are:

- DSP Windowing
- Spectral Interpolation

Each of these techniques will be discussed in more detail.

Spectral Interpolation

Spectral interpolation is a technique that was proposed1 as a means of controlling interference in a coherent sine wave measurement. Briefly, this techniques looks at the interference components on either side of the desired signal in the frequency domain. These nearby components are used to estimate, through interpolation, the interference component at the signal frequency. The interpolation estimate is then subtracted from the observed value at the signal frequency.

Several claims are made for spectral interpolation. These include:

- eliminates spreading the signal into multiple frequency components
- eliminates time-domain distortion of measured signal
- is intuitively more straight forward than windowing

While these claims may be true, they do not relate to the objectives of any sine-wave measurement. In particular, these claims do not relate to the accuracy, repeatability, or cost of a particular technique.

The actual spectral interpolation technique is very simple. The basic hypothesis is that the contribution of interfering signals on the value of a particular signal component can be predicted by interpolating an interference value from the interference in nearby bins that are adjacent to the signal bin. For example, a first order estimate of the signal component minus the interpolated interference is:

signal - $\text{interference}_{est}$ = -.5*bin(n-1) + bin(n) -.5*bin(n+1)

where signal, $\text{interference}_{est}$, and bin(i) represent the following complex amplitudes:

Signal	The measured complex amplitude of signal (including interference) at analysis freq.
$\text{Interference}_{est}$	The estimate of interference based on the amplitude of nearby frequencies
Bin (i)	The measured complex amplitude of the i^{th} component of Fourier transform

The following table shows the weighting of nearby bins for the proposed first and second order spectral interpolations:

Order	**Weighting of spectral component**				
bin	**n-2**	**n-1**	**n**	**n+1**	**n+2**
First	0	-.5	1	-.5	0
Second	.1666667	-.6666667	1	-.6666667	.1666667

Windowing

Windowing is a traditional approach to analyzing non-periodic signals using periodic signal analysis tools. Most DSP measurement algorithms are based on the assumption that the data set being analyzed contains a whole number of cycles of a periodic signal that repeats forever. Most real world signals (except for carefully constructed test signals) do not satisfy this constraint. When non-periodic signals are analyzed as if they are periodic, the last sample of the collected data set is assumed to be followed by the first, second, etc. samples of the same sequence. This usually introduces a significant step discontinuity. As indicated earlier, this discontinuity results in signal components unrelated to the signal being analyzed. Windowing techniques weight each sample in the collected data sequence by appropriate amounts to minimize the effect of the discontinuity.

Windows and windowing techniques have been discussed extensively in the literature2 3 4. Also, there has been unpublished usage of windowing to limit the effects of transients5. For this paper several different windows will be considered. These include:

- Hanning Window
- Blackmann Window
- Rosenfeld Window
- Optimum 4-term window

Each of these windows is constructed by adding a DC term to cosine components that are harmonics of the measurement period. In general, a DSP window can be written in the form:

$w(t)=a_0 + a_1\cos(\omega t) + a_2\cos(2\omega t) + a_3\cos(3\omega t)$ . . .

over the interval t=0 to t=T and ω equals 2π/T. The following table defines each of the windows in terms of their coefficients, a_n:

Window	a_0	a_1	a_2	a_3
Hanning	0.5	-0.5		
Blackman	0.42	-0.5	0.08	
Rosenfeld	0.762	-1.0	0.238	
Opt. 4th	0.3635819	-0.4891775	0.1365995	-0.0106511

Any of these DSP windows can be applied to a measured sample sequence by creating a DSP window of the selected type that contains a number of samples equal to the number of measured samples. Then each sample in the measured sample sequence is multiplied or weighted by the corresponding sample in the DSP window. Since these windows all approach zero at the beginning and end of the window sequence, they all reduce the magnitude of the difference between the first and last sample, and therefore the interference at signal frequency produced by under settled transients.

Each of these windows has different characteristics when used to control the effect of noncoherent interference. Their effects will be discussed in more detail later.

The signal spectrum of a windowed signal can be related to the signal spectrum of the same signal without windowing. In particular, the value of each Fourier component of the windowed component can be expressed as a sum of the primary component plus nearby components. (This is a consequence of the Fourier transform equivalence which says that multiplying two signals in the time domain is equivalent to convolving their Fourier transforms).

For example, if a signal contains a sine-wave component that falls exactly in bin ‘n’ is windowed, then the signal component in bin ‘n’ of the windowed signal can be expressed as a weighted sum of components from the Fourier spectrum of the unwindowed signal. The following table shows how the content in the signal bin of the windowed signal is obtained from a weighted sum of the Fourier spectral bins near bin n in the unwindowed spectrum:

Window	**Weighting of unwindowed spectral component**						
bin	**n-3**	**n-2**	**n-1**	**n**	**n+1**	**n+2**	**n+3**
Hanning	0	0	-.5	1	-.5	0	0
Blackman	0	0 .095238	-.59523	1	-.59523	.095238	0
Rosenfeld	0	.156167	-.65617	1	-.65617	.156167	0
Opt. 4th	-.01465	.187968	-.67313	1	-.67313	.187968	-.01465

It is important to note that the mathematical effect of multiplying by a Hanning window is identical to the effect of first order spectral interpolation. Also, the effect of the Rosenfeld window is similar but not exactly the same as second order spectral estimation.

Controlling the Effects of Transients

Since Spectral Interpolation and DSP windowing are simply different techniques for achieving the same result (i.e. creating a weighted sum of a signal bin with nearby bins), the following analysis results are presented for windowing with the understanding that they apply equally to the corresponding spectral interpolation.

It was assumed that the signal being measured consist of an exact sinewave plus a transient component. The transient component is usually composed of one or more components which are either purely damped exponentials or damped sinusoids. The analysis of a variety of different transients, including both exponentials and low-frequency damped sinusoids produced a very interesting result. While the sinewave components of the transient at low analysis frequencies depended upon the specific shape of the transient, the high frequency components depended only on the discontinuity (difference) between the sample that would have been collected after the last sample in the measured sequence, compared with the first sample in measured sequence.

The interference cause by unsettled transients is a result of adding the signal-frequency component of the transient to the signal frequency component that would be obtained if there were no transient present. The upper curve in Figure 2 shows the error introduced by components of a transient with a one volt discontinuity at various signal frequencies when measuring a one-volt sinewave (assuming worst case phase relationships, see below). If the signal is windowed using a Hanning window prior to analysis, error introduced by the same one volt discontinuity on a one volt sinewave at various signal frequencies is modified as shown by the lower curve in Figure 2.

The exact effect of the transient interference will depend on the exact phase of the sinewave signal compared with the phase of the interference component at the same frequency. If the two components are exactly 90 degrees out of phase, the interference will have little effect on the observed signal amplitude. However, if the interference is exactly in phase are 180 degrees out of phase, its amplitude will either add or subtract directly from the amplitude of the desired sinewave. Figure 2 shows the error introduced in a one volt sinewave measurement made in the presence of a transient with a one volt discontinuity. Note that if the ratio of the transient discontinuity to signal amplitude is different(not equal to one), then the error (in dB) will be multiplied by that ratio. For example, if the discontinuity caused by the transient is five volts, the error introduced into the signal measurement will be five times as large as the errors shown in Figure 2.

Comparison Criteria

There are three basic criteria for comparing techniques for minimizing the effect of noncoherent components on sine-wave measurements. These are:

- Computational complexity or cost
- Noise gain
- In-band residual

Computational complexity is proportional to the amount of computation that is needed to implement a given technique. For this discussion, the number of multiplies required to perform a given computation will be used as a measure of computational complexity. For example: computing a Fourier transform using one of the fast algorithms may offer a factor of 100 improvement over a direct implementation.

Noise gain is the apparent increase in in-band noise as a result of a particular technique. For example: a technique that adds up the content of five bins to arrive at a signal level adds up five bins of noise energy as well. If the total amplitude of the measured quantity is not increased, the noise will increase by the square root of five or 7 dB.

In-band residual is the amount of extraneous signal that remains after applying the technique.

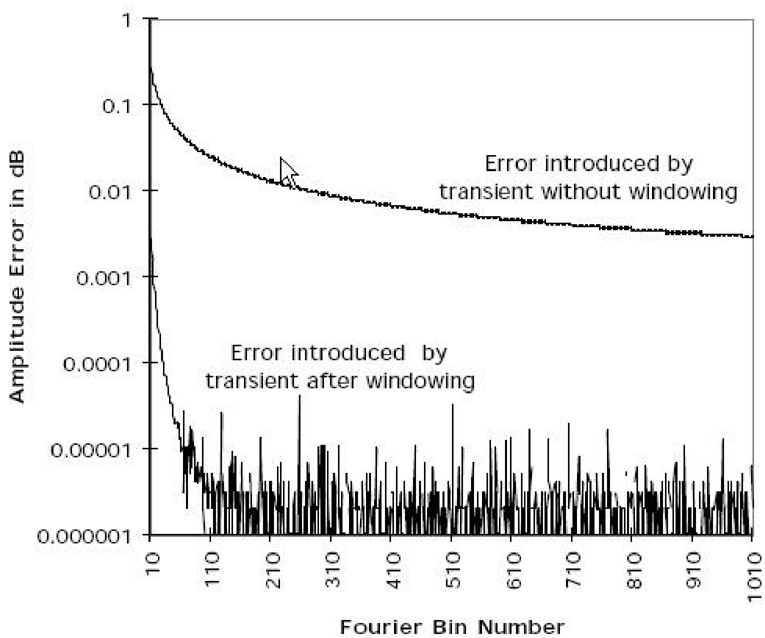

Figure 2 - Error introduced in measuring one volt sinewave by one volt transient, with and without windowing

While computational complexity is directly related to the cost of applying the technique, noise gain and in-band residue relate to the accuracy and repeatability of the measurement. Any noise or residue combines with the desired signal through vector addition. As a result, noise or residue components can lead to apparent changes in signal amplitude or phase. As the noise and residue are decreased, the error introduced in the measurement is also reduced.

Windowing vs. Spectral Estimation

The only effective way of comparing these techniques is by comparing their computational complexity, noise gain, and in-band residue. Each of these will be considered in turn.

Computational Complexity

There are two different implementations of both the windowing and spectral estimation technique. One is based on direct computation of individual spectral lines. The other is based on using FFT techniques. Where a small number of frequency components need to be computed, direct computation is often faster than an FFT. However, where many frequency components are needed the FFT will be much faster.

Using direct computation of frequency components, the component at each frequency of interest is computed by multiplying each element by the corresponding elements of sine and cosine sequences at the analysis frequency. Direct computation requires 2N multiplies and additions per frequency component. By comparison, an FFT (order 2) requires 2N multiplies and additions per stage of the process. The number of stages is equal to the $\log_2 N$. The total number of multiplies will be at least $2N\log_2 N$. An FFT may also have additional computation time associated with coefficient computation and reordering the result (FFTs frequently scramble the order in which samples appear) . Thus a 1024 point FFT with 10 stages will have approximately 20,000 multiplies (with precomputed coefficients).

Based on the above estimates, for 1024 samples, it is usually more efficient to use an FFT algorithm, as opposed to direct computation, when more than 10 signal components need to be computed.

To perform a harmonic distortion test using direct computation of the fundamental, 2nd and 3rd harmonic would require on the order of 6N multiplies and additions. When windowing is added, an additional N multiplies will be performed. Thus the number of multiplies required to window and compute values is 7N. If FFTs are used to perform the same computation (window + harmonics), $(2\log_2 N+1)N$ multiplies are needed. For 1024 samples, $(2\log_2 N+1)N$ is equal to 21N. Clearly, a direct computation is significantly more efficient.

To perform harmonic distortion using spectral interpolation, at least three analysis frequencies must be computed to obtain a value for each signal component. Thus, a total of nine direct computations are required, resulting in a computational complexity of 18N multiplies. If an FFT is used, the amount of complexity is essentially equal to that of an FFT. Thus it would be $2N\log_2 N$. As indicated above, an FFT requires 20N multiplies when N is equal to 1024. Thus for 1024 samples there will be little difference which method of computation is used. For sample sequences of 512 points or less, it will probably be faster to use an FFT. However, for longer sequences, direct computation will be more efficient.

The following table compares the number of multiplies for each harmonic distortion analysis technique, for both short (where an FFT is better) and long sequences (where direct computation is better):

Step	Windowed Direct	Windowed FFT	S.I. Direct	S.I. FFT
DSP Window	N	N	Not Used	Not Used
Compute	6N	$2N\log_2 N$	18N	$N\log_2 N$
Frequency components				
Total Multiples	7N	$N (2\log_2 N+1)$	18N	$2N\log_2 N$
N=8	56	56	144	48
N=16	112	144	288	128
N=32	224	352	576	320
N=64	448	832	1152	768
N=128	896	1920	2304	1792
N=256	1792	4352	4608	4096
N=512	3584	9728	9216	9216
N=1024	7168	21504	18432	20480
Range of preferred N	N≥8	N≤8	N≥512	N≤512

An examination of this table reveals that for virtually all practical cases (cases where the number of signal samples is greater than sixteen, the practical lower limit for measuring third harmonic distortion) using direct computation of windowed data will be significantly faster than using an FFT and/or spectral interpolation. For typical sequence lengths of 128 to 1024 samples, the advantage will be a factor of 2 to 3 improvement in favor of the windowed technique.

Noise Gain

For windowed measurements, the noise gain will be equal to the peak-to-RMS ratio of the window used. The table below gives the noise gain for several common windows.

For the spectral interpolation technique, the noise gain will be computed by applying the square root of the sum of the squares of the individual terms that are combined to compute a corrected measurement. For example, in the case of first order interpolation, the corrected value is computed using the following formula:

corrected value for bin n=bin n-(bin (n+1) + bin(n-1))/2

Assuming the RMS value of the noise in each bin is v_n , the total noise in the corrected value will be v_n multiplied by the sqrt(1+ $1/2^2$ + $1/2^2$) or sqrt(1.5) or 1.22 (or 1.76dB). This means that the random noise in the signal bin will be 1.76 dB higher than it would be if spectral interpolation were not used. The amount of noise increase will depend on the order of the interpolation used. The table below also shows the noise gain for first order interpolation.

Technique	**Noise Gain in dB**
Hamming Window	1.76
Blackman Window	2.37
Rosenfeld Window	2.81
Opt. 4th Window	2.96
First Order S.I.	1.76

While none of the techniques imposes a severe penalty, using the Hanning window or the first order spectral interpolation will produce the best (smallest) noise gain. This is further confirmation of the mathematical equivalence of the Hanning window vs. first order spectral interpolation.

In-band Residual

Is stated earlier, windowing and the comparable spectral interpolation are identical in their effect on reducing the residual error from an unsettled transient. In this regard the techniques are identical.

Practical Benefit

DSP windowing has been employed at LTX for several years as a means of reducing settling time for testing Subscriber Line Circuits (SLIC). For one device, the overall test time without windowing was approximately 12 seconds. With windowing, the test time was reduced to approximately 2.5 seconds, without loss of measurement quality.

Conclusion

The use of windowing techniques can serve a useful purpose with coherent measurements to reduce the effect of unsettled transients at the output of a DUT. Windowing used for this purpose can enable significant test time reduction.

In addition, the technique called “spectral interpolation” was found to be mathematically equivalent to windowing. However, its computation was found to be significantly more time consuming for the combined measurement of fundamental, second and third harmonic amplitudes.

Based on these results, it can reasonably be expected that windowing is the best method of reducing the effects of unsettled transients.

[1] Mark Burns, “Improving DSP-Based Measurements with Spectral Interpolation”, Proceedings of the International Test Conference, 1995, pp355

[2] Lawrence R. Rabiner and Bernard Gold, “Theory and applications of digital signal processing”, Prentice-Hall,1975

[3] Eric Rosenfeld, “DSP Measurement of Frequency”, Proceedings of the 1986 International Test Conference, pp981

[4] Jin Choi, “Understanding Windows in DSP”, unpublished (presented at 1994 LTX Corporation Synchro Users Conference)

[5] Takahashi-San, LTX Corporation, Tokyo, Japan (unpublished)