A Computationally Efficient Method for Accurate Spectral Testing without Requiring Coherent Sampling

Zhongjun Yu, Degang Chen, Randy Geiger
Iowa State University

Abstract

The Fast Fourier Transform is the ubiquitous method of choice for spectral testing. However, its correct application to periodic signals requires either strict coherent sampling, or careful windowing, or other techniques that are computationally inefficient. This paper introduces a new method for achieving accurate spectral testing for periodic signals without the need for coherent sampling or windowing. Furthermore the method is computationally very efficient with only minimal addition to the computational complexity of FFT. The method is validated with both simulation data and experimental data. Extensive controlled simulation indicates that the method is very robust to errors in signal frequency, phase, amplitude, additive noise, and so on. Statistical analysis and comparative studies demonstrate that the proposed method achieves spectral testing accuracies similar to those obtained with perfect coherent sampling in an ideal noise-free environment.

1. Introduction

In many important application areas such as signal processing and communications, spectral performance of an integrated circuit is of critical concern. The DFT (Discrete Fourier Transform) [1], or the fast implementation of it, FFT, is the most prevalent method for spectral performance testing. However, when using the FFT for spectral testing of periodic signals, one must be extremely careful not to allow the so-called frequency leakage problem to distort the actual spectrum of the signal. Specifically, one must make sure that the data record being used in the FFT algorithm represents exactly an integer number of periods of the signal. In other words, the signal frequency and the sampling clock frequency of the data acquisition system must be exactly coherent with each other. The FFT algorithm is notorious for being extremely intolerant to even the slightest mismatches between the two frequencies. It is shown that frequency errors at the levels of a small fractional ppm can cause disastrous measurement results. The resultant error manifests itself as the frequency leakage phenomenon in which energy from the fundamental spectral line is spread into neighboring frequencies causing the appearance of a “skirt” around the spectral line.

Figure 1 shows the incorrect spectra of four pure sine wave signals. The details are irrelevant but notice that the spectra are qualitatively distorted due to straightforward application of FFT with non-coherently sampled data sets. A correct spectrum should consist of only two spectral lines at the input signal frequencies.

When a periodic signal is not a pure sine wave, its distortion components cause spectral lines to appear at integer multiples of the fundamental frequency. The heights of these harmonic spectral lines relative to the fundamental are key specifications in spectral testing of a signal. However, when the skirt due to non-coherency becomes higher than the harmonic distortion spectral lines, the spectral testing results will be erroneous. To combat the leakage or skirting problem, the IEEE standard [2] as well as industry best practice is to require coherent sampling, meaning that the clock signal of the data acquisition system should be perfectly synchronized with the signal under test so that an integer multiple of signal periods are captured in a data record of length M. When this is guaranteed, direct use of FFT is permitted and the data analysis is computationally very efficient, requiring only O(MlogM) operations.

A second method is to use the windowing technique [3] while allowing non-coherent sampling. This technique does not remove the skirting due to non-coherency; rather it merely suppresses the skirting levels at frequencies far away from the base frequency. By doing so it alters the heights of the original spectral lines. Care must be taken in order to correctly recover the spectral lines. Another limitation is due to the fact that the amount of skirt suppression is limited and hence it is difficult to apply to relatively high purity signals.

Other methods for combating spectral leakage include singular value decomposition [4], in which the singular values of an M X M matrix formed from the data record are computed with time complexity $O(M^3)$, 2-D FFT [5], which requires > $O(M^2 \log^2 M)$ operations, and filter banks [6]. These methods are accurate but they are computationally very inefficient.

In this paper, we introduce a new method for achieving very accurate spectral testing for periodic signals without the need for coherent sampling or windowing. The key idea can be called fundamental identification and replacement. The method is targeted for high precision spectral testing and is very efficient computationally with O(MlogM) operations. In its current form, it is limited to signals that are close to being sinusoidal. In the next section, we will reformulate the DFT problem for periodic signals and point out the leakage mechanism. In section 3, we present the proposed method in detail. Sections 4 and 5 contain simulation and experimental results validating the proposed method. It is shown that the method is very robust to errors in signal frequency, phase, amplitude, additive noise, and so on. Section 6 presents statistical analysis of extensive simulation results showing that the method achieves spectral testing accuracies comparable to those obtained with perfect coherent sampling in an ideal noise-free environment.

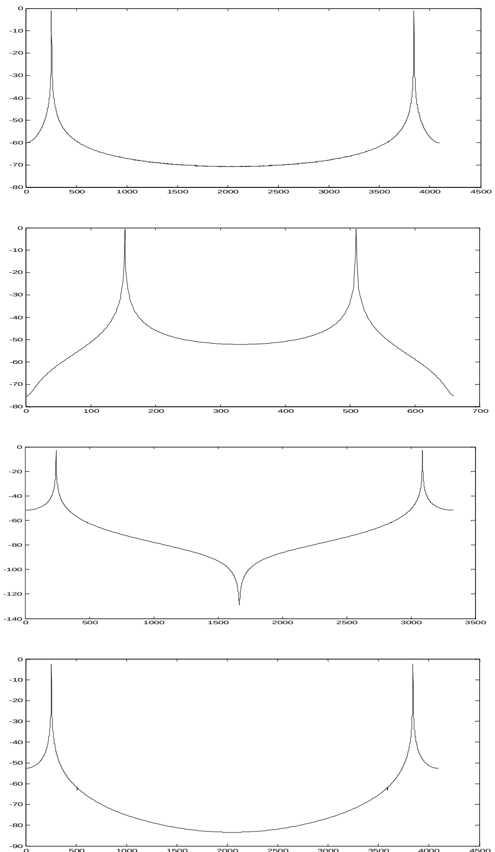

Figure 1. Spectral leakage due to non-coherency

2. DFT of Periodic Signals

Let f_s be the sampling frequency, $T_s = 1/f_s$ the sampling interval, f_i the unknown input signal frequency, and M_0 the nominal data record length. Then $J = M_0 f_i / f_s = J_0 + \Delta$ will be the number of periods of input signal in the data record, where J_0 is the integer part of *J*, Δ is the fraction part of *J*. J_0 and M_0 are assumed to be co-prime. Δ is unknown, so is *J* (J_0 could be known).

Let the input signal be:

$$x(t) = A\sin(2\pi f_i t + \theta) + \sum_{n\ge 2}(b_n \sin 2\pi n f_i t + a_n \cos 2\pi n f_i t)$$

where $A \approx 1$, $\theta \in [0, 2\pi)$, a_n, b_n for $n \ge 2$ are all unknown, but together they satisfy $\sum_{n\ge 2}(b_n^2 + a_n^2) << 1$,

$$\left|\sum_{n\ge 2}(b_n \sin 2\pi n f_i t + a_n \cos 2\pi n f_i t)\right| << 1.$$

The samples of $x(t)$ at sampling rate f_s are given by:

$$x[k] = A\sin(2\pi f_i \frac{k}{f_s} + \theta) + \sum_{n\ge 2}(b_n \sin 2\pi n f_i \frac{k}{f_s} + a_n \cos 2\pi n f_i \frac{k}{f_s})$$

Since $\frac{f_i}{f_s} = \frac{J}{M_0} = \frac{J_0 + \Delta}{M_0} = \frac{J_0}{M_0} + \frac{\Delta}{M_0}$, we also have

$$x[k] = A\sin(2\pi \frac{J_0}{M_0} k + 2\pi \frac{\Delta}{M_0} k + \theta)$$
$$+ \sum_{n\ge 2}(b_n \sin 2\pi n f_i \frac{k}{f_s} + a_n \cos 2\pi n f_i \frac{k}{f_s})$$
$$= \frac{A}{2j}(e^{j\theta} e^{j2\pi \frac{\Delta}{M_0} k} e^{j2\pi \frac{J_0}{M_0} k} - e^{-j\theta} e^{-j2\pi \frac{\Delta}{M_0} k} e^{-j2\pi \frac{J_0}{M_0} k})$$
$$+ \sum_{n\ge 2}(b_n \sin 2\pi n f_i \frac{k}{f_s} + a_n \cos 2\pi n f_i \frac{k}{f_s})$$
$$= x_1[k] + x_h[k]$$

where $x_1[k]$ is the base harmonic component of $x[k]$, $x_h[k]$ is the sum of 2nd and higher harmonic components of $x[k]$. Taking the Fourier transform leads to:

$$X(f) = F[x](f) = F[x_1](f) + F[x_h](f)$$
$$= X_1(f) + X_h(f)$$

Using the DFT formula $X[n] = \frac{1}{M_0}\sum_{k=0}^{M_0-1} x[k] e^{-j\frac{2\pi}{M_0}nk}$,

X_1[n] can be derived to be:

$$X_1[n] = \frac{A}{2j}[e^{j\theta}\delta(n - J_0) - e^{-j\theta}\delta(n + J_0)]$$
$$+ \frac{A}{2jM_0} e^{j\theta} \sum_{k=0}^{M_0-1}(e^{j\frac{2\pi}{M_0}\Delta k} - 1) e^{j\frac{2\pi}{M_0}(J_0 k - nk)}$$
$$- \frac{A}{2jM_0} e^{-j\theta} \sum_{k=0}^{M_0-1}(e^{-j\frac{2\pi}{M_0}\Delta k} - 1) e^{-j\frac{2\pi}{M_0}(nk + J_0 k)}$$

In the above equation, the term in the first line corresponds to the correct spectral line of the Fourier transform of the periodic signal x_1(t), the terms in the second and third lines will be non-zero as long as Δ is non-zero. These terms give rise to a skirt around the main spectral line. They are a linear combination of M_0 different frequency terms and therefore have M_0 independent basis functions. This is the reason that the skirt due to non-coherency can exhibit so many qualitatively different shapes. This very fact also dictates that any attempt trying to identify the skirt in the frequency domain with a reduced number of basis functions is going to be ineffectual.

On the other hand, if one can somehow efficiently identify θ and Δ, then the skirt term can be calculated, and subtracted from $F[x]$ so that

$$\hat{F}[x] = F[x] - skirt \;\; = F[x_1 + F[x_h] - skirt$$

$$= idealterm + skirt + F[x_h] - skirt \;\; = idealterm + F[x_h]$$

Therefore, $F[x_h]$ is recovered, which is the harmonic distortion component in $x[k]$. Even in such a case, the problem remains that there is no fast algorithm to compute the "skirt" terms. The computation of the skirt terms in $X_1[n]$ involves, for each n:

$2M_0 + 1$ exponential evaluations

$3M_0 + a\ few more$ multiplications

$2M_0 + 1$ additions

Multiplying the above by M_0 for $n = 0,1,2,\cdots M_0 - 1$, leads to a total computation $O(M_0^2) >> M_0 \log M_0 >> M_0$.

In the next section we introduce a time domain identification and replacement method for removing the skirt but using only O(M_0logM_0) operations.

3. The Proposed Method

From the discussion in the previous section, we know that as long as Δ is non-zero, which means the data record length is not exactly an integer number of signal periods, the DFT algorithm introduces an error term in the Fourier transform of the fundamental component. This leakage term can be so large that it completely inundates the harmonic distortion components, making it impossible to correctly test the true spectrum of the signal. We also mentioned that if one can remove the skirt term, then the distortion terms can be revealed and correctly tested. Our goal is then to find a method for estimating and removing the skirt term from the DFT spectrum. If this can be done efficiently, the harmonic distortion components will show as spectral lines in the spectrum. Then the spectral heights at the harmonic frequencies can be correctly computed to determine the signal's spectral performance.

Instead of identifying and removing the skirt in the frequency domain, the proposed approach works indirectly in the time domain. From the captured data (with distortion), we first estimate the amplitude, frequency and phase of the fundamental harmonic component. Once this is done, we replace the non-coherent first harmonic component with a sine component that has the same amplitude and phase but a slightly modified frequency so that it becomes coherent with the sampling clock. This is all done in the time domain and it will be followed by standard FFT spectral analysis.

The enabling underline premise is that harmonic distortion components are all at frequencies that are multiples of the first harmonic. This means that they are all orthogonal to the component that we are trying to identify and hence we can ensure that with the right algorithm they will have minimal effects on the accuracy of the first harmonic identification. Furthermore, the harmonic distortion components are assumed to be much smaller than the base harmonic.

It is important to point out that in the process of first harmonic identification and replacement, the higher order harmonic distortion components are unchanged. On one hand this is good since distortion computation is not affected. On the other hand, this also means that any skirting effects in the harmonic distortion components are not corrected since these components are non-coherent with the sampling clock either. The result is that the measured heights of harmonic distortion components may be reduced by their own leakage effect. Fortunately, as is well known, the leakage only causes the spectral line to be lower by a small fraction of one dB. Therefore, its effect can be ignored. In the proposed method, we also select an appropriate data record length so that the data is close to being coherent to begin with. This further assures that the skirting effects on the distortion components can be comfortably neglected.

As denoted in the previous section, let the number of integer periods in the data record be J_0 and the total number of periods be $J_0 + \Delta = J$. Then the base harmonic in the signal can be represented by

$$x_1[k] = A_0 \sin(2\pi \frac{J_0 + \Delta}{M_0} k + \theta)$$

There are three unknown parameters A_0, Δ, and θ. Therefore a minimum of three known data points is needed to identify them. For example,

$$k = 0: \qquad x_1[0] = A_0 \sin(\theta)$$

$$k = M_0: \qquad x_1[M_0] = A_0 \sin(2\pi\Delta + \theta)$$

$$k = \frac{M_0}{2}: \qquad x_1[\frac{M_0}{2}] = A_0 \sin(\frac{2\pi J_0}{2} + \pi\Delta + \theta)$$

contain sufficient information to identify the three parameters. However, there are several difficulties. First, these equations are nonlinear and nontrivial to solve. Second, the values of the base harmonic component $x_1(t)$ are unknown. Only the distorted total signal x(t) is measured and known. Finally, measurement noise and quantization error in the measured data will affect the accuracy of the parameter identification.

We now introduce two different methods for dealing with these difficulties. In the first method we make use of two facts: 1) the signal purity is high with total distortion energy in the < −60 dB to −80 dB range, 2) the distortion components are going to be measured relative to the first harmonic component.

Fact 2 means we can scale the acquired data so that A_0 = 1 without affecting spectral testing. We can then use fact 1 to help achieve the correct scaling. Specifically, we scale the measured data record (with distortion) so that it has total signal power equal to 0.5. Since the total signal power is equal to the base harmonic power plus the total harmonic distortion power, we have

$$0.5 = 0.5A_0^2 + TDP$$

$$A_0 = (1-2TDP)^{0.5} \approx 1 - TDP \approx 1$$

where TDP stands for the total distortion power which is assumed to be in the −60 dB to −80 dB range. Hence the error in taking A_0=1 is in the 0.01 to 1 ppm range and negligible. Furthermore, the follow up computation of Δ and θ is robust with respect to small errors in A_0. All together, this can lead to a small error in the first harmonic replacement. If the error magnitude is at the −50 dB level or lower, its contribution to the spectrum will be at the −100 dB level or lower. Such error levels will likely not affect the spectral testing results.

By taking $A_0 = 1$, the above equation becomes

$$\therefore x[0] \approx x_1[0] \approx \sin(\theta)$$

$$x[M_0] \approx x_1[M_0] \approx \sin(2\pi\Delta + \theta)$$

$$\therefore 2\pi\Delta + \theta = \sin^{-1}(x[M_0])$$

or:

$$\theta = \sin^{-1}(x[0])$$

$$2\pi\Delta = \sin^{-1}(x[M_0]) - \sin^{-1}(x[0])$$

To further improve the robustness of the identification results, we will use more data points to obtain redundant solutions and use averaging to reduce the effects of noise and quantization. For example, we can use the following equations

$$x[1] = x_1[1] \approx \sin(\frac{2\pi}{M_0}(J_0 + \Delta) + \theta)$$

$$\therefore \frac{2\pi}{M_0}J_0 + \frac{2\pi}{M_0}\Delta + \theta = \sin^{-1}(x[1])$$

$$x[M_0 - 1] \approx \sin(\frac{2\pi}{M_0}(J_0 + \Delta)(M_0 - 1) + \theta)$$

$$\therefore \frac{2\pi}{M_0}(J_0 + \Delta)(M_0 - 1) + \theta = \sin^{-1}(x[M_0 - 1])$$

to obtain another set of solutions for Δ and θ.

Once A_0, θ, Δ are computed, we can generate a new data record by replacing the base harmonic from the original data (which is sampled non-coherently and causes possibly large skirts) with one that is coherent with the sampling clock. This is done by simply subtracting a sine component with the identified parameters and adding a sine component with the same A_0 and θ but with Δ being rounded to zero.

$$\hat{x}[n] = x[n] - \hat{A}_0 \sin(\frac{2\pi(J_0 + \hat{\Delta})}{M_0}n + \hat{\theta}) + \hat{A}_0 \sin(\frac{2\pi J_0}{M_0}n + \hat{\theta})$$

$$= \hat{A}_0 \sin(\frac{2\pi J_0}{M_0}n + \hat{\theta}) + \sum original harmonics$$

$$+ \{A_0 \sin(\frac{2\pi(J_0 + \Delta)}{M_0}n + \theta) - \hat{A}_0 \sin(\frac{2\pi(J_0 + \hat{\Delta})}{M_0}n + \hat{\theta})\}$$

The first term on the right hand side is a coherent sine wave, its DFT will have spectral lines with magnitude $\hat{A}_0$ at the $(1+J_0)th$ and $(M_0 - J_0)th$ frequency bins. No skirts will come from the first term. The second term represents the harmonic distortion components in the original data records. In order for the harmonic distortion components to correctly show up in the DFT of $\hat{x}[n]$, the third term magnitude needs to be sufficiently below the expected harmonic distortion level. As we commented above, the power-based normalization can help us reduce the power level of the third term to the −100 dB or lower level, assuming the total harmonic distortion is at the −60 dB level.

Additional strategies for reducing the error:

1. Use more equations than necessary to estimate A_0, θ, Δ, and use least square method in the estimation. The error effects due to approximating $x[n]$ by $\hat{x}[n]$ are likely to be uncorrelated with each other. The least square method has the capability of removing uncorrelated equation errors from the solutions. Hence the estimation errors can be reduced to below the harmonic distortion level.
2. In selecting the equation to use for solving A_0, θ, Δ, care can be taken so that the harmonic distortion effects can be minimized. For example, avoiding using data points that are at the same phase angle for the second harmonic frequency can reduce the effects of the 2nd harmonic distortion on A_0, θ, Δ estimation. Similarly, effects of other harmonic distortion terms can be reduced.
3. Use zero-crossing up-edge trigger to start the data acquisition. By doing so, θ will be approximately 0 and the errors in estimation θ will have less effects.
4. Choose a data record length to be possibly different from M_0 so as to minimize Δ. For example, take $2M_0$ samples instead of M_0 samples. Search through samples M_0 to $2M_0$, to find the $(1+M_1)th$ data point that most closely matches the 1st point in the data sequence. That is, $x[1]$ through $x[M_1+1]$ most closely match an integer number of signal periods. Then use $x[1]$ to $x[M_1]$ for spectral analysis. This method will lead to a small Δ. In most cases, the fact that Δ is very small will be sufficient to reduce the skirt to a level that will allow the correct computation of harmonic distortion terms without the need of

replacing the non-coherent first harmonic component. This will be seen from the examples.

The first method can be summarized by the following steps:

1. Synchronize the digitizer and input signal by using a positive zero-crossing edge to trigger sampling
2. Take a sufficient number of samples (e.g. $2M_0$ instead of the regular M_0 samples)
3. Find a data point between M_0 and $2M_0$ that best matches the first data point and use all the points before this as the data record
4. Identify base harmonic by first normalizing the data power to be 0.5 (and taking A_0=1) and then using earlier equations to compute Δ and θ
5. Replace the non-coherent base harmonic by a coherent base harmonic as on last page
6. Perform FFT analysis as usual

In the second method for identifying the three base harmonic parameters, the frequency or Δ is identified first. Then it is used with a least square method to identify the Cartesian form of A_0 and θ, that is, $A_0 \sin\theta$ and $A_0 \cos\theta$. This method requires slightly more computation and is most suitable for situations where normalization is not appropriate. For the sake of space, we will simply describe the following step by step procedure for using the method.

1. Take a sufficient number of samples (maybe a few times the intended FFT length, no need for synchronization or edge triggering)
2. Search among a subset of the data points (eg the 10% points that are closest to zero-crossing) for a few best-matched pairs
3. Select the pair (call them $x[k_1]$ and $x[k_2]$) whose index difference ($k_2 - k_1$) has the most number of factors and use all the data point between this pair together with one of the end points (eg $x[k_1]$ to $x[k_2-1]$) as the data record
4. Count the integer number of cycles J_0 in the data record and computer the fractional cycle
$$\Delta = \frac{1}{2\pi}\sin^{-1}\left(\frac{x[k_2]-x[k_1]}{\sqrt{\hat{A}_0^2 - x^2[k_1]}}\right)$$
where $\hat{A}_0$ is a first estimate of the base harmonic magnitude. Then the input signal frequency is
$$f_{in} = f_s \frac{J_0 + \Delta}{k_2 - k_1}$$
5. At a subset of data points write
$$x[k] = A_0 \cos(\theta)\sin(2\pi f_{in} t_k) + A_0 \sin(\theta)\cos(2\pi f_{in} t_k)$$
and use least square method to identify $A_0 \cos(\theta)$ and $A_0 \sin(\theta)$
6. Perform the first harmonic replacement
$$\hat{x}[k] = x[k] - A_0 \sin(2\pi f_{in} t_k + \theta) + A_0 \sin(2\pi f_s t_k J_0 /(k_2 - k_1) + \theta)$$
7. Perform FFT analysis as usual

Both of the two methods presented above are computationally very efficient. In both cases, data acquisition time will be longer than if perfect coherent sampling is available. The time can be up to a factor of 5 depending on how many extra data points are to be collected. In any case, the data acquisition time should be small for no more than 10K points with today's ADCs. In the first algorithm, the search for the best match and the normalization both require computational time that is proportional to the data record length. Computation required for Δ and θ is very small. In the second method, steps 2, 4 and 6 each require a time proportional to the data record length. Step 5 is more flexible but it can also require a time proportional to the data length if we want to take the most advantage of the noise reduction power of least squares. The computation in step 3 is trivial but the results of step 3 can affect the time requirement of step 7, the standard FFT. If the data record length has small factors, then the FFT is very efficient and requires O(MlogM) operations where M is the length. In the unlikely event when M is a prime number, the FFT would require $O(M^2)$ operations. Because of this, the first method could be slower if it ends up using a prime number FFT length. Nevertheless, in both method 1 and method 2, the most dominant part of the computation is for performing the standard FFT. Therefore both methods are extremely efficient.

4. Simulation results

To verify the performance of the proposed algorithms, we have conducted extensive simulation study as well as experimental study. In this section we will present two spectral testing examples, one represents a typical case and the other a less-likely case. The simulation environment is set up so that many parameters are randomly generated. The intended data record length is randomly chosen among $2^{10 \sim 13}$. The quantization resolution of the digitizer is randomly selected among 12 ~ 17 bits. The input signal frequency is generated by selecting a random ratio of f_{sig} to f_{samp}. The signal magnitude has random error, but guaranteed to be within a range. A random DC off set error is introduced to the signal. A random phase synchronization error is also included. Additive measurement noise is introduced at the input node of the digitizer with a standard deviation of around 1 LSB. The distorted sine wave signal is generated by adding random amount of harmonic distortion components to a pure sine wave. For comparison, four different spectral testing methods are simulated. These are: 1) perfect coherent sampling with ideal noise-free environment, 2) straightforward application of DFT assuming periodic sampled sequence, 3) simply selecting a best data record length as discussed in the proposed method, and 4) the proposed method which includes 3) followed by first harmonic identification and replacement.

First example

This is a representative example in that 1) the ideal case and proposed method generated a spectrum showing zero or minimal skirts as can be seen from Figure 3 and Figure 6, 2) these two methods and the 3rd method all produced accurate measurements of the signal SFDR, as can be seen from Figures 3, 5, and 6, 3) straightforward application of DFT suffered from large errors due to non-coherency, as is easily seen from Figure 4. The true SFDR of the signal is computed analytically in continuous time domain. A time domain illustration of the coherent and non-coherent data samples is shown in Figure 2. It is particularly worth pointing out that the proposed method (Figure 6) is subject to all sorts of non-idealities whereas in Figure 3 everything is ideal.

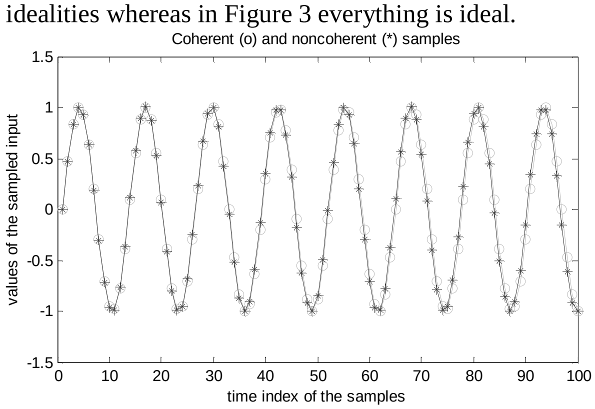

Figure2. Time domain data samples from coherent and non-coherent sampling

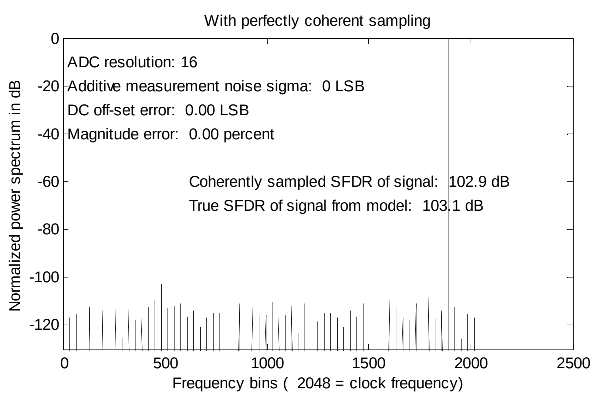

Figure 3. Spectrum from ideal coherent sampling

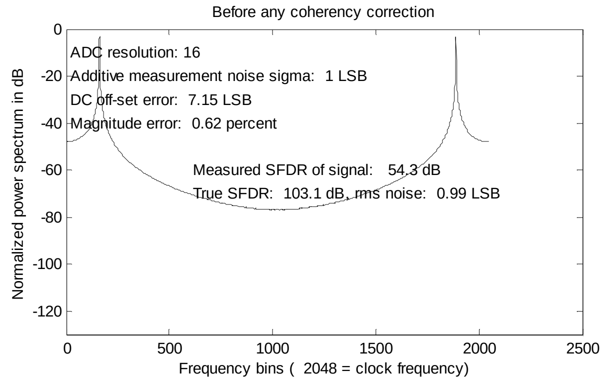

Figure 4 Straightforward application of DFT to non-coherent data samples

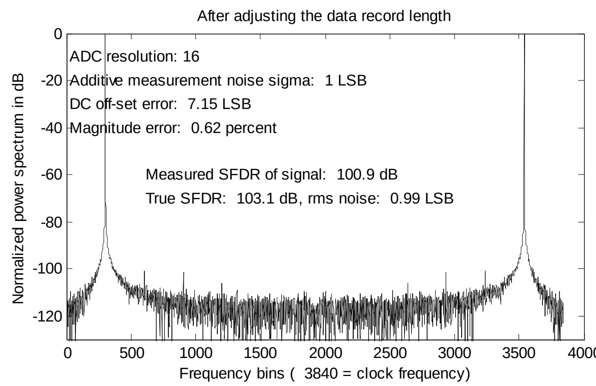

Figure 5. DFT with non-coherent samples with best data record length

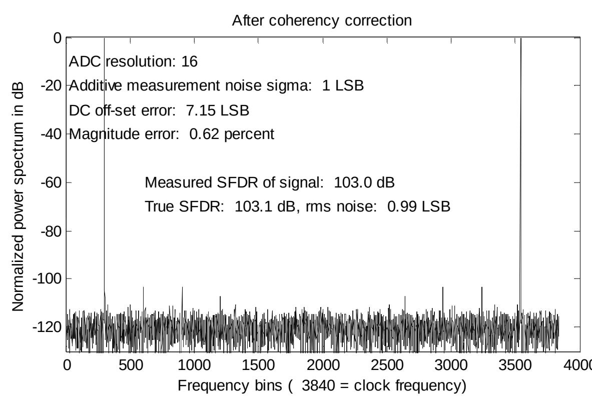

Figure 6. The proposed method: non-coherent samples, best data length, and 1st harmonic replacement

Second example

From our experience of large numbers of simulations, the situation illustrated in the first example is the most commonly seen. This means in many situations the signal spectral distortion can be computed to a sufficient accuracy by following part of the procedures described in the proposed method so that an appropriate data record length is selected to approximately match the coherent sampling condition.

In such cases, the skirting is not removed from the spectrum and therefore the signal spectrum is not correctly computed, but the error due to skirting is not large enough to dominate the harmonic distortion components. However, with a relatively small but certainly not small enough to be negligible probability, a situation can happen that the un-corrected skirting is still too large for the correct testing of harmonic distortion in the signal. The second example illustrates this situation. In this example the signal frequency is relatively high as compared to the Nyquist frequency. Hence the signal phase changes very quickly with the index. This makes it more difficult to find a data record length that is very close to being coherent. Figure 7 illustrate the time domain data samples from coherent and non-coherent sampling. As is from the previous example, the straight forward application of FFT to the non-coherent samples produced a spectrum that is totally dominated by the spectral leakage effect and the signal distortion cannot be correctly measured, see Figure 8. By following partially the procedures of the proposed method to select an approximately coherent data record length, the skirting effect due to non-coherency is greatly reduced as can be seen from Figure 9. However, the skirt is still above the harmonic distortion component level and also significantly above the noise floor of the 15-bit ADC. The spectrum in deed looks quite different from the correct spectrum in Figure 10 which is obtained using perfect coherent sampling in an ideal testing environment. Fortunately, the proposed method with base harmonic identification and replacement is still capable of removing the skirting effect and producing a spectrum (see Figure 11) that is essentially the same as that in Figure 10. Again, the proposed method worked with non-coherent data samples and the testing environment is subject to various non-idealities. The computation requirement is only slightly larger. This clearly demonstrates both the efficacy and the robustness of the proposed method.

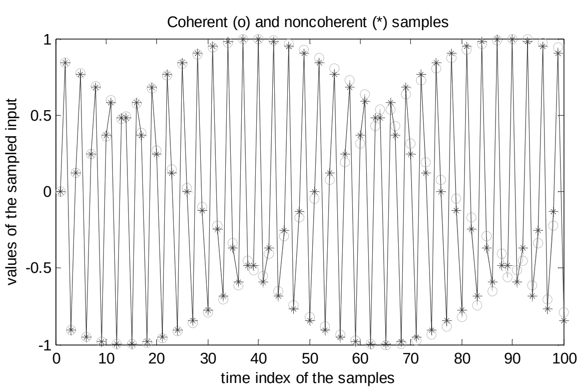

Figure 7. Data samples from coherent and non-coherent sampling

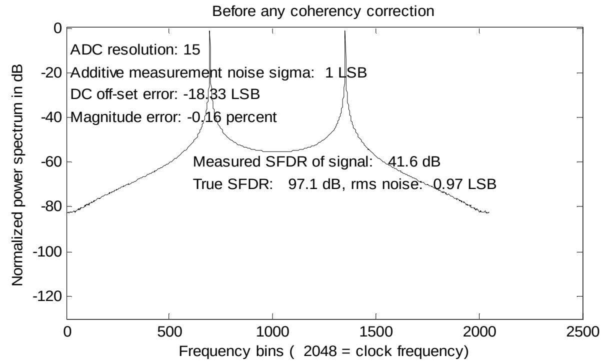

Figure 8 Straightforward application of DFT to the non-coherent data samples

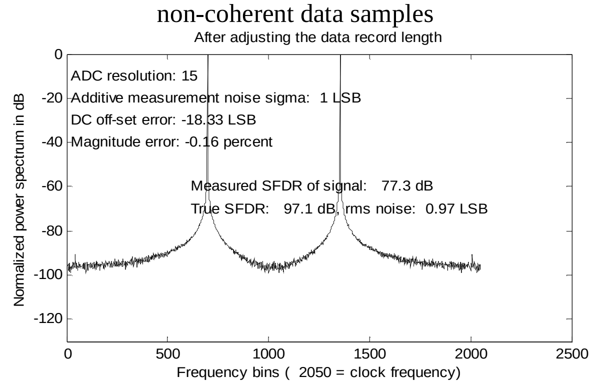

Figure 9. DFT with non-coherent samples and best data record length without first harmonic replacement

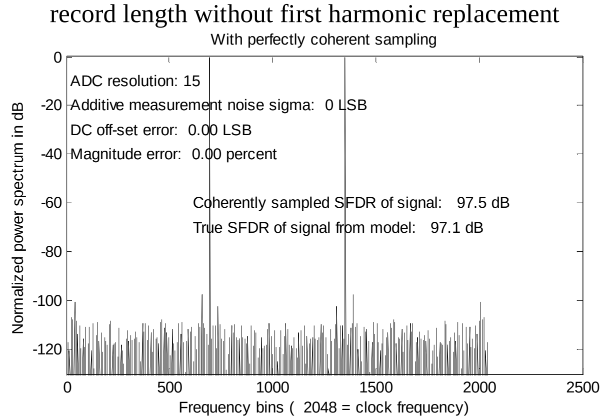

Figure 10. Signal spectrum from perfect coherent sampling and ideal test environment

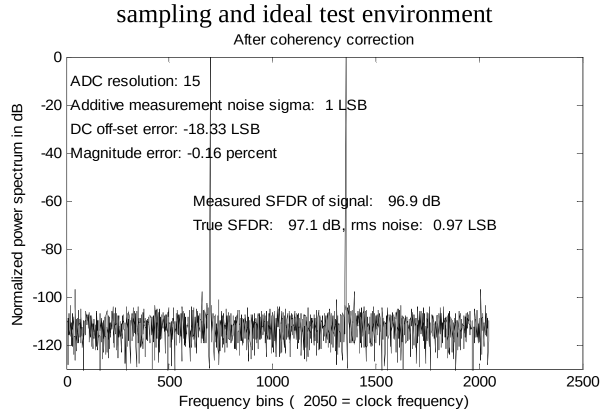

Figure 11. The proposed method: non-coherent samples, best data length, and 1st harmonic replacement

5. Experimental Results

Since the proposed method exhibited excellent spectral performance with sufficient robustness to various error sources, we wanted to validate the algorithms with experimental data. Figure 12 is the time domain representation of the captured data. The data is collected in an industry setting. Although we are not in a position to talk about details of how the data is collected but it suffices to say that it is supposed to be discarded. The data was sampled non-coherently and the signal was slightly clipped at both the top and the bottom. A small zoomed-in piece is show in Figure 13 in which one can see the clipping at the top.

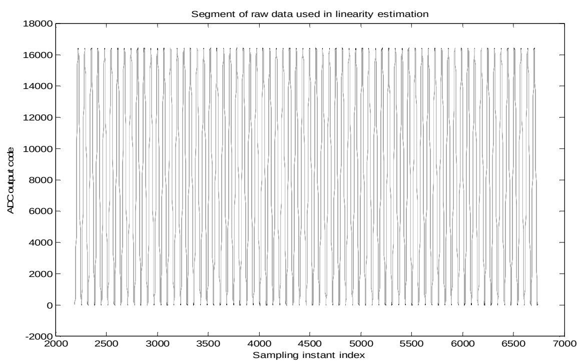

Figure 12 Raw data of non-coherent samples of sine wave slightly clipped at the top and bottom

To analyze the spectral contents of the signal, one can straight forwardly apply DFT to the raw data. The resultant spectrum is shown in Figure 14. The clearly visible skirting completely dominates any harmonic distortion components that may be present. Hence no testing information is obtained.

Next the proposed method was applied to analyze the spectral contents of the captured signal. The resultant spectrum is shown in Figure 15. Notice that all skirting effects have been removed and the noise floor has been pushed down to the −100dB level. Rich spectral contents are clearly shown.

Since the signal has not been coherently sampled, we do not know what the true spectrum of the signal is. However we do know that the signal is a nice clean sine wave clipped at the top and bottom. By adjusting the amount of clipping, we can create a clipped sine wave whose true spectrum is shown in Figure 16. It exhibits very similar spectral characteristics to those in Figure 15. Therefore we believe that the spectrum shown in Figure 15 is indeed the true spectrum of the signal being captured.

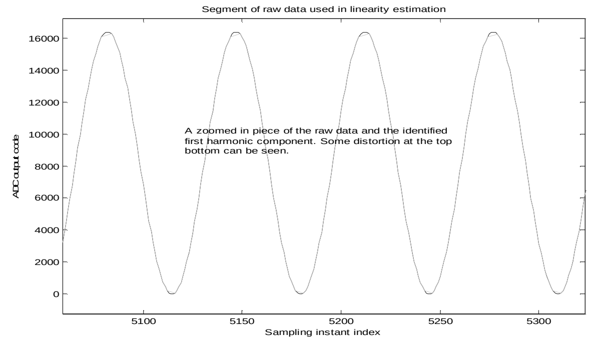

Figure 13. A zoomed in piece showing clipping at the extremes.

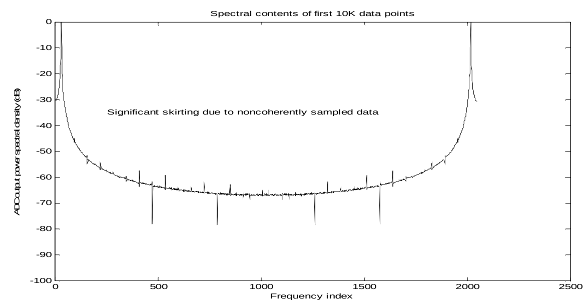

Figure 14. Spectrum by standard application of DFT

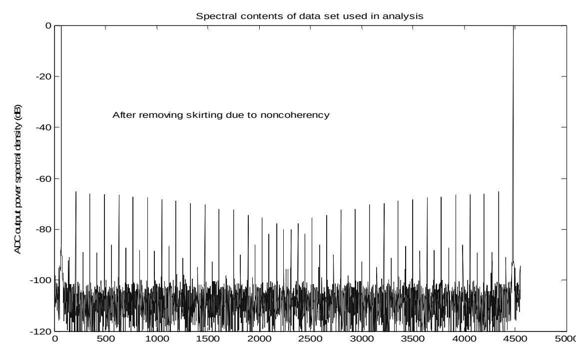

Figure 15. Spectrum of the captured signal using the proposed method

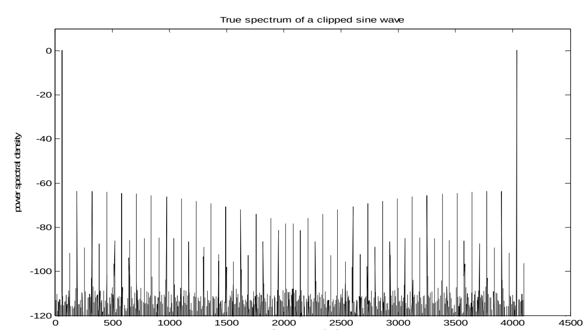

Figure 16. True spectrum of a sine wave with clipping at the top and bottom

In addition to the spectral information about the signal, Figure 15 also reveals some information about the linearity of the underlining digitizer. Notice that the harmonic distortion spectral lines due to clipping have a nice smooth envelope resembling a sinc function. The highest harmonics spectral line that deviates from this nice envelope is around the −85dB level. This deviation is due to the addition of distortion from the digitizer and it means that the digitizer is likely to have a SFDR near the 85dB level. In deed the digitizer used to capture the data is a 14 bit ADC. The use of the proposed method in testing ADC spectral performance is currently under study.

6. Statistical Performance Study

In the simulation results section, we have mentioned that the simulation program is set up so that many non-ideal factors are randomly generated and many runs with different randomly generated parameters have been tested. There is one phenomenon, though not frequent, that caught our attention. That is, once in a great while very large SFDR testing error is produced with the proposed method. We decided to perform some extensive statistical analysis to figure out what the cause is. Well over 10000 runs were conducted with various randomly generated parameters such as input signal frequency, digitizer resolution, the size of the fractional period in the data record, total data record length, the amount of signal magnitude errors, initial synchronization errors, additive measurement noise, and so on. Correlation between the signal SFDR testing errors and various controlled parameters was studied. We also studied the correlation to the accuracy in A_0, Δ, and θ estimation. Because of page limitation, we cannot discuss all of our findings here. But it is sufficient to point out that the proposed method is very insensitive to all the parameters that we studied except for one: the ratio of the input signal frequency to the sampling clock frequency. Figure 17 contains data from over 10000 runs. The horizontal axis is the ratio of input signal frequency to clock frequency as a percentage. The vertical axis is the signal SFDR testing error from each run. Clearly, when the signal frequency is very close or above the digitizer's Nyquist frequency, the SFDR measurement errors are usually large. There are also certain signal frequencies, for example exactly 1/3 or ¼ of the clock frequency, for which the proposed method does not work. In hindsight, the reason is simple. In such cases, one or more of the harmonic frequencies or their aliased versions become exactly equal to some other harmonic frequencies. When this happens, correct measurements cannot be obtained. We also want to point out that at these frequencies, the perfect coherent sampling with ideal environment cannot guarantee correct measurement of the signal spectral performance either.

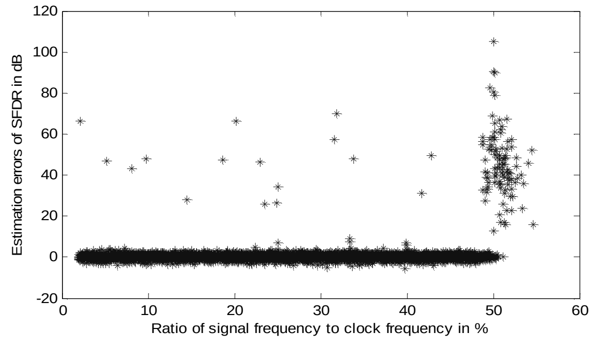

Figure 17 Correlation of SFDR estimation error to signal frequency from over 10K random runs

Based on the findings from the statistical study, we recommend that certain signal and clock frequency combinations should be avoided if the ratio f_{in}/f_s can be reduced to a rational number with small integers. Next, we coded this requirement into the code where random input frequencies are generated so that no such frequencies are generated. The generated frequency is rounded to a nearest coherent frequency for use in the benchmark ideal case for comparison. We also avoided using extremely low frequencies, since it is known that at very low frequencies parameter identification accuracy is difficult to maintain. [7]

Simulation of 1000 cases that avoid inappropriate frequencies was conducted. Figure 18 illustrates the signal SFDR testing errors when the proposed method is used. Again the horizontal axis is the input signal frequency to clock frequency ratio. Notice that in all 1000 runs, the SFDR errors are within ±4dB. For comparison, Figure 19 illustrates the SFDR testing errors for the same 1000 cases when the ideal perfect coherent sampling is used. For this method, there is one case when the frequency was rounded, it became a degenerate frequency of exactly ¼ of clock frequency. The testing result for this case was totally off. There are two other cases when the frequencies were rounded, they became those frequencies that should have been avoided. Other than these three cases, the ideal perfect coherent sampling method performed just as well as or slightly better than the proposed method. Figure 20 illustrates the results when DFT is straight forwardly applied. It can be clearly concluded that such a method cannot be used for non-coherently sampled data. Figure 21 illustrates the results by partially following the procedures described in this paper to select the most suitable data record length but without first harmonic replacement. It can be seen that for a large percentage of the cases, this method produced sufficiently small testing errors.

Table 1 summarizes the statistics of these comparative study results.

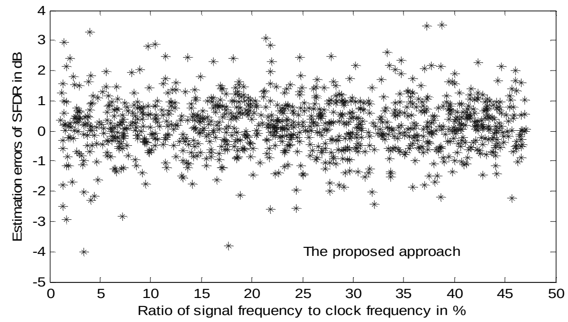

Figure 18. SFDR testing errors in 1000 runs using the proposed method, vs f_{in}/f_s*100

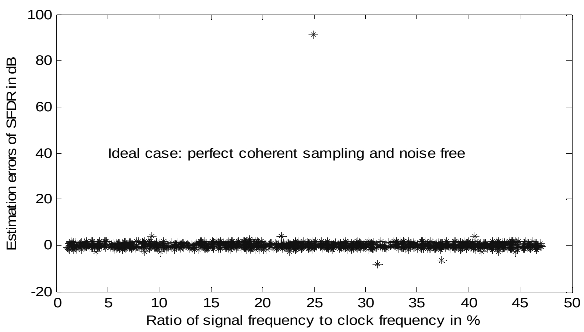

Figure 19. SFDR testing errors in 1000 runs using the perfect coherent sampling in the ideal case

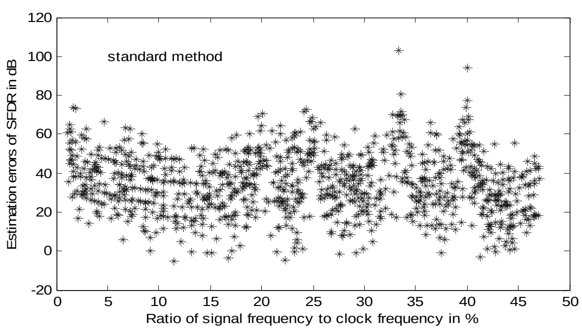

Figure 20. SFDR testing errors in 1000 runs non-coherent sampling in the standard approach

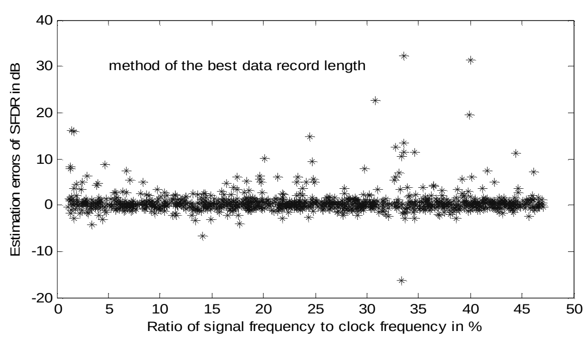

Figure 21. SFDR testing errors in 1000 runs using non-coherent samples with best data length

Table 1 Summary of SFDR test errors in 1000 runs

	Ideal	Standard	Best length	Proposed
min	-8.33	-5.26	-16.35	-4.01
max	91.1	103.2	32.19	3.52
mean	0.01	34.75	0.61	0.17
std	3.08	15.80	2.63	0.91

*Ideal case was skewed by three bad data points, would have been comparable to proposed case if those bad points were excluded

Conclusion

A new method for accurate spectral performance testing is presented that does not require coherent sampling or the use of windowing. The proposed method achieves SFDR testing performance comparable to ideal perfect coherent sampling. Extensive simulation study and statistical performance analysis have been conducted. The method is demonstrated to be very robust to various sources of errors. The computational efficiency of the algorithm is excellent with only slightly more computations than FFT. However, the method is vulnerable when the signal frequency is near the Nyquist frequency or at certain degenerate frequencies for which even DFT with ideal perfect coherent sampling cannot work well. Therefore, the proposed method offers comparable performance to FFT with perfect coherent sampling but without requiring coherent sampling.

References

1. A. Oppenheim, *et al*, *Discrete-time Signal Processing*, Prentice-Hall, 1999.
2. IEEE Std 1057-1994 (R2001), “IEEE Standard for Digitizing Waveform Recorders.”
3. P. Carbone, *et al*, “Windows for ADC Dynamic Testing via Frequency-Domain Analysis,” *IEEE Trans. Instr. & Meas.*, 50(6), pp 1571-1576, 2001.
4. J. Zhang ands.J. Ovaska, “ADC characterization based on singular value decomposition,” *IEEE Trans. Instr. & Meas.*, 51(1), 2002.
5. X. Gao, *et al*, “Analysis of second-order harmonic distortion of ADC using bispectrum,” *IEEE Trans. Instr. & Meas.*, 45(1), pp. 50-55, 1996.
6. C. Rebai, *et al*, “Non-coheret Spectral Analysis of ADC Using Filter Bank,” *IEEE Instr. & Meas. Tech. Conf.*, pp. 183-187, 2002.
7. S.M. Kay, *Fundamental of Statistical Signal Processing: Estimation Theory*, Prentice-Hall, 1993.