REPRINT

Title: **THE INTERPOLATION OF SPARSE TIME HISTORY DATA**
Author: **Strether Smith**
Source: **Lockheed Palo Alto Research Laboratory**

Improvements in signal aliasing protection in modern digital data acquisition systems allow the collection of discrete data at sample rates that are very close to the theoretical limit of two times the desired data bandwidth. This produces very high "efficiency" in the acquisition process and very good fidelity and near-optimum bandwidth for spectral measurements. However, the resulting time history data is very sparse and must be processed to reconstruct the waveform.

A number of techniques have been described in the literature to calculate the intermediate values. One, that can be implemented by many of the data processors available, is to use zero-insertion in the spectral domain, a method that works perfectly for signal segments that do not have significant energy at the Nyquist frequency. Unfortunately, most signals that are to be interpolated do not satisfy this condition and significant errors result.

This paper describes an extension to the zero-insertion technique that uses "windowed guard bands" to reduce the errors to acceptable levels. The interrelationship between sample ratio, interpolation ratio, window used, and guard-band size is examined and their effect on errors is discussed. A procedure that assures errors of less than .1% of full scale is described.

DISCUSSION If data acquired by a digital data acquisition system is to be used directly to represent the waveform, conventional wisdom (and a variety of sources, [1],[2]) dictates that at least ten points per cycle of the highest frequency component of interest should be acquired. This sample ratio[1] produces a data representation like that shown in Figure 1. As can be seen, the waveform is fairly well defined and that the peak value can be determined to an accuracy of 5%.

On the other hand, Shannon's sampling theorem states that, if a data set is band-limited, it can be completely described if it is sampled at a rate that is "slightly more" than twice the frequency of the highest-frequency component in the data. The problem is that it must be guaranteed that there are no significant signal components above one half of the sample rate (the Nyquist frequency) and this condition must be assured by the combination of sample rate and anti-alias filters [3]. Modern data acquisition systems have anti-aliasing filter systems that provide data that is band limited to a level of 80 dB or better. In particular, converters based on oversampling and digital filtering ("sigma-delta") provide over 90 dB of aliasing protection when sampling ratios as low as 2.2 are used. Thus, data can be acquired with sampling ratios of much less than ten with insignificant aliasing errors.

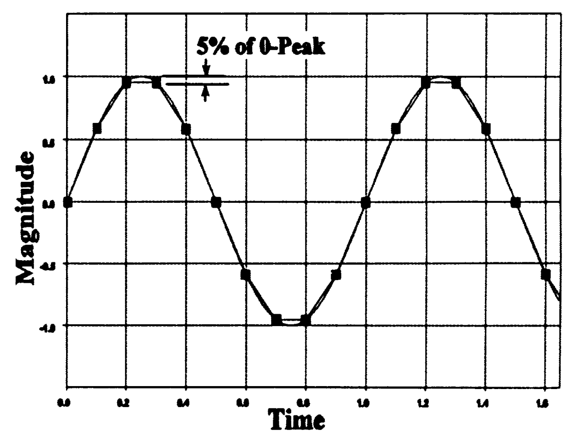

Figure 1. Ten Points/Cycle Define a Waveform Reasonably Well.

The use of low sampling ratios has enormous advantages in that both the acquisition speed and the amount of data stored are accordingly reduced. If three points per cycle (Figure 2), rather than the "conventional-wisdom" value of ten, are acquired then the hardware requirements are reduced by 70%.

As can be seen, the data acquired at low sample ratios is very sparse and that the raw data is not a good representation of the real data set. The problem is how to take advantage of the premise of Shannon's theorem and reconstruct the data between the acquired points. If we can do this, we can both reduce our hardware requirements and produce a more complete (accurate?) data set.

[1] Sample Ratio is defined as the sample rate divided by the maximum frequency of interest.

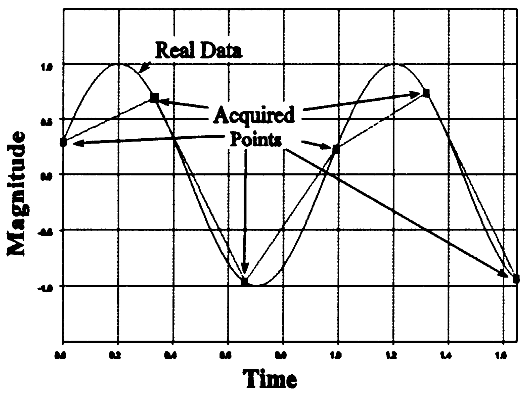

Figure 2. Three Points/Cycle is Too Sparse to Define the Waveform Directly.

The basic mathematics of data interpolation have been well established for years. The approaches, which may be applied in either the time domain or the spectral domain, are based on the fundamental concept of Fourier/harmonic analysis that any finite-length, frequency-limited, signal can be expressed as a superposition of sine and cosine terms. For this discussion, we will use the "spectral-domain approach" because it is straightforward in concept and can be implemented in a variety of commercially-available time-series-analysis data processors[2].

The basic principles of time-domain interpolation in the spectral domain are described in [4]. To interpolate by a factor of K (calculate K – 1 intermediate values between acquired points) the procedure is (Figure 3):

1. Select the time range to be interpolated (N points = 8).
2. Calculate the spectrum of the selected range via Fourier transform (N complex spectral values). Divide the Nyquist-Frequency value by 2 and assign each half to a separate component (N + 1 values).
3. Insert [(N – 1) × K – 1] zeroes between the Nyquist components to produce N × K complex spectral values.
4. Use the inverse Fourier transform to calculate the interpolated time history. (K × N points).

For the example shown, where there is an integer number of cycles in the selected buffer, this procedure works perfectly. However, success is contingent on the fact that the portion of the signal to be interpolated is made up of signals that have an integer number of waves in the buffer. This is not the general, or even the usual, case.

[2] This study was performed using DADiSP on a PC-Clone computer.

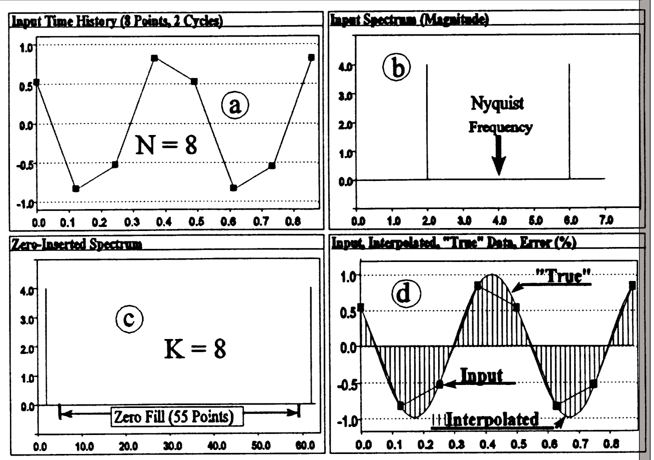

Figure 3. Zero-Insertion Interpolation With Integer Waves per Buffer. a) Data Points Selected for Interpo-

Figure 4 shows the same procedure applied to an example of the general case where the selected data set is not made up of "integer-waves-per-buffer" sets. The impact of this is:

* The spectrum of the selected data is "spread." The critical effect for this application is that the spectral magnitude is not "small" near the Nyquist frequency.
* When the zero-insertion is performed, there is a discontinuity in the spectrum at the Nyquist frequency.
* When the inverse transform is performed, ringing (error) is found at the beginning and end of the selected data window. (Actually, the errors extend all the way through the data set but are largest at the ends).

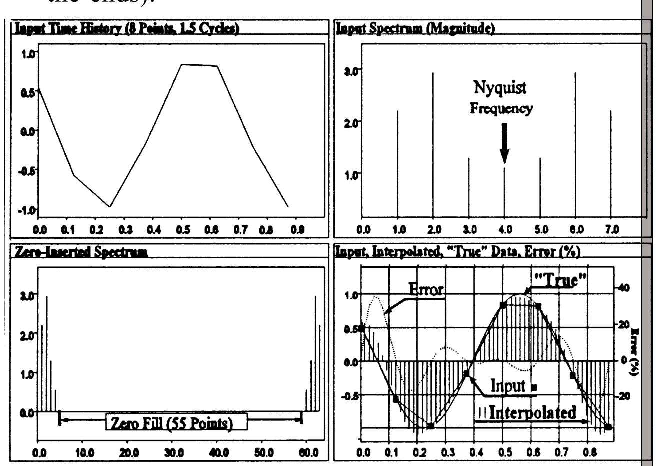

Figure 4. Zero-Insertion Interpolation with a Non-Integer (1.5) Waves/Buffer Results in Significant Errors.

The problem arises because of the fundamental assumption of harmonic analysis that the time window being analyzed is repeating. For the example data, the analysis procedure "thinks" that the data looks like Figure 5. The ringing that is seen at the boundaries is the result of the discontinuities. The only time that this will not occur is when there is an integer number of waves in each buffer so that there is no discontinuity at the ends.

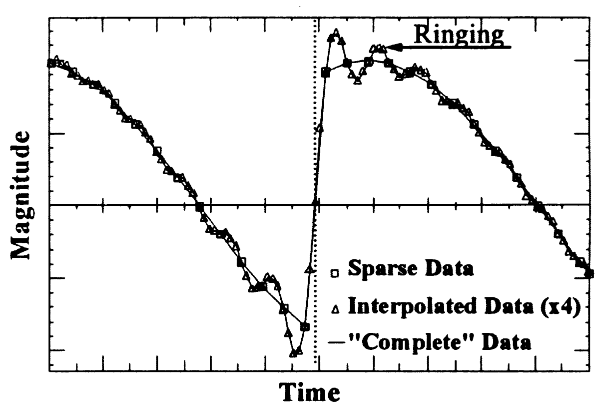

Figure 5. The "False" Discontinuity at Buffer Ends Causes Errors in Interpolation.

Fourier analysis provides us with a tool that allows us to quantify the size of the ringing. In the time domain, a discontinuity is propagated through the data in accordance with the following equation.

$\sin(\pi x)/(\pi x) = \sinh(\pi x)$

x = spacing (in samples) between the discontinuity and the point of interest.

Thus, the errors that occur for a "false" discontinuity are:

Zero at the initially-acquired points

$\sin(\pi x) = 0$ for x = 1, 2, 3.....

Non-Zero for points between the initially-acquired points where we want to interpolate ($x \neq$ integer). The error in the calculation for the interpolated data point half way between the first and second data points away from the discontinuity (x = 0.5) can be as large as 49% of full scale, depending on the character of the discontinuity.

This consideration gives us a clue to a solution. Since the effect of the buffer-end mismatch is attenuated as we get away from the discontinuity, a possible approach might be to include a larger data set in our interpolation calculation and extract the desired interpolation range from the larger result. This approach is called "guard banding" [5] which means to include data on either side of our range of interest to act as a buffer so that the transients are reduced to acceptable levels. This procedure is diagrammed in Figure 6.

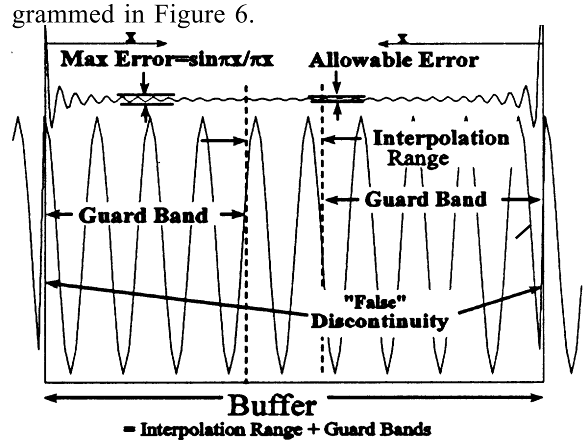

Figure 6. Guard Bands May be Used to Isolate Discontinuities.

The size of the required guard band may be directly calculated from the SINH function. The maximum potential error due to a unit discontinuity is equal to $1/\pi n$ where n is the number of (original) data points away from the discontinuity. Thus, if a guaranteed accuracy of 0.1% of full scale is necessary, then the false discontinuity must be more than $n = 1/0.001\pi = 318$ points from our desired range. Thus we must perform calculations on 636+ range-of-interest points to achieve the desired results.

This approach has two notable drawbacks:

* We cannot perform interpolation any closer than 318 points from the end of our data set.
* An inordinately-large number of calculations must be performed.

To help with these deficiencies, we resort to windowing of our expanded data set. The objective of the windowing is to reduce the effect of the discontinuity so that a smaller guard band may be used.

The choice of the best windowing function is not obvious. In exchange for reducing the effects of buffer-end discontinuities, windows also broaden the effective filters in the Fourier analysis and this will cause distortion.

Preliminary studies were performed using the following windows.

* Triangle (Bartlett)
* Sine
* Sine with flat section in the interpolation range.
* Three-term Blackman (Figure 7).
  $0.42-0.5\times\cos[2\times N/(L-1)] + .08\times\cos[4\times N/(L-1)]$

N = Point in Buffer $0 \le N \le L - 1$

L = Buffer Length

This window was chosen because of its rapid sidelobe attenuation to -60 dB and its time-history attenuation characteristics near the end of the buffer.

A few experiments showed that the Blackman window provided the best results.

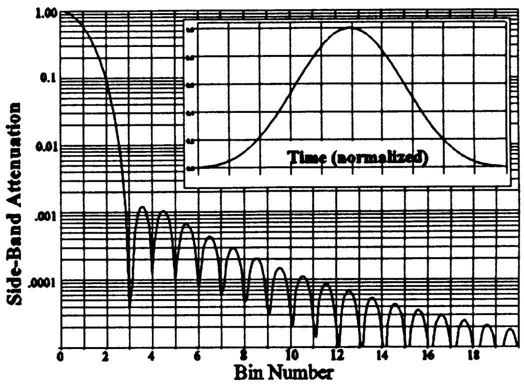

Figure 7. 3-point Blackman window, Time and Spectrum Shape

To see the effect of the windowing, we will use data set similar to the one in Figure 2. It was constructed with the following parameters:

* Sample Rate = 8 sample/second.
* Frequency = 1.5 Hz (5.33 points/cycle).
* Interpolation range of 8 points (1 second, 1.5 Waves/Buffer).

Figure 8 shows the results when a 10-point guard-band and Blackman window are used. The errors are reduced to well below the desired 0.1% level from the 40% level found when no guard bands were used.

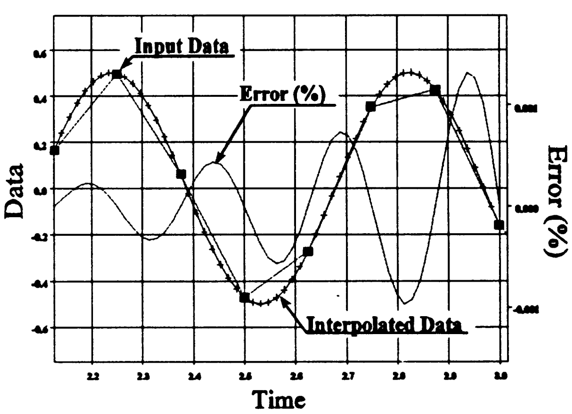

Figure 8. Data set with 1.5 Cycles/Buffer Interpolated With Windowed Guard Bands.

Once the Blackman window was selected, the size of the guard band required to achieve a desired level of error was determined. An empirical study was performed with the following parameters.

* The objective was to find a window/guard-band length that produced an interpolated time history with a given magnitude of ***peak*** error. A peak error level of 0.1% of full scale was chosen as the acceptability criterion.
* The variables covered were:
  * Interpolation Range= 4, 16, 256, 512 (points).
  * Interpolation Ratio = 8 and 16.
* Two input time history types were used as test cases:
  * Cosine time histories that were designed to include full-scale discontinuities at the ends of the padded buffer. The frequencies were chosen to provide sampling ratios from 2.2 to 3.37.
  * Gaussian-random time histories were used as a sanity check to demonstrate that the cosine history produced the worst results.

Figure 9 shows a typical result set. The error (% of full scale) is plotted against guard band length for a 16 point set that was interpolated by 8. It can be seen that the worst errors are incurred at low sample ratios and that a guard band of 30 or more points assures that errors for any sample ratio greater than 2.2 will be less than 0.1%.

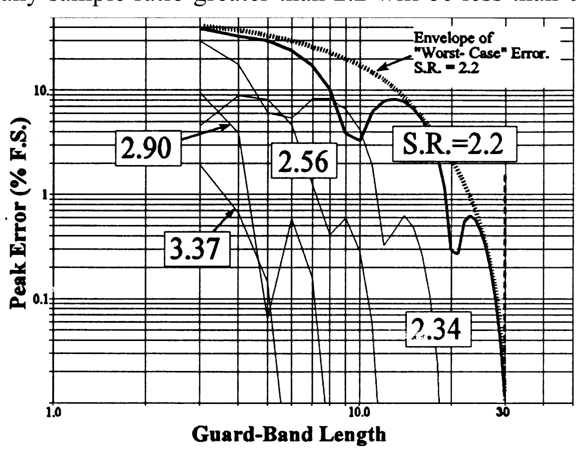

Figure 9. Peak Error vs. Guard-Band Length for Several Sample Ratios. Interpolation Range = 16 Points, Interpolation Ratio = 8.

Figure 10 shows the error envelopes for data sets with 2.2 points/cycle interpolated with a variety of different numbers of input points and interpolation ratios. The error curves are enveloped to indicate the worst case for "short" and "long" interpolation ranges. As can be seen, a guard band of less than 30 points produces errors of less than 0.1% for all cases.

For longer input time histories, it is recommended that a guard band that is proportional to the input length be used to avoid numerical problems when dividing by the window shape. If a rule of not using the window where the multiplier is less than .05 then the minimum guard-band length should be the interpolation length/8 for the Blackman window. Thus, the guard band length should be:

For interpolation ranges of up to 240 points:

30 points

For more than 240 points:

Number of Points/8

Thus, if the input interpolation range is 100 points, the guard band is 30 and the total data set included is 160 points. Because the "power-of-2-length" Fourier Transform is so efficient, it is probably appropriate to include 256 points. This will improve both the interpolation accuracy and calculation speed.

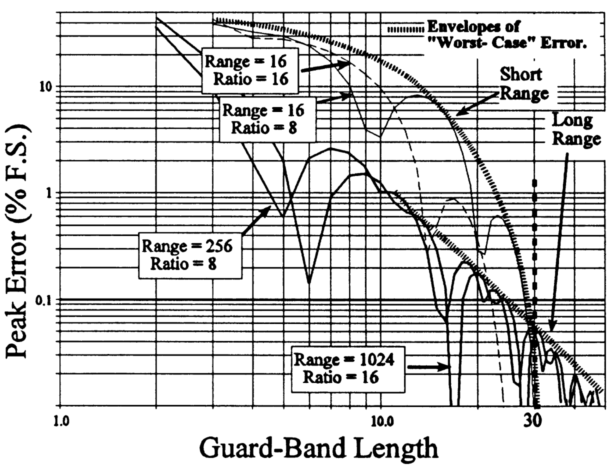

Figure 10. Calculated Peak Error and Envelope vs. Guard Band Length for a Variety of Interpolation Ranges and Ratios. Cosine Input, Sample Ratio = 2.2

IMPLEMENTATION

The interpolation process is diagrammed in Figure 11. The steps are:

1. Select the range of data **(n points)** to be interpolated from the full, acquired, data set.
2. Determine the size of the guard band **(G points)**, based on the number of points selected, using the criteria in the previous section.
   a. Determine the size of the complete set **(N = n+2 × G)**.
   b. Optional: Expand the set to the next higher power of 2 (N = 32, 64,....1024..) so that the Fourier Transform processing will be more efficient.
3. Retrieve the selected data set.
4. Multiply the data set by the Blackman Window.
5. Calculate the (complex, two-sided) spectrum with the Fourier Transform.
6. Split the spectrum at the Nyquist Frequency. The Nyquist frequency component must have special handling. In most FFT algorithms the component is divided by two and each half is assigned to the upper and lower portions of the spectrum. The result is N + 1 spectral values.
7. Insert **(K – 1) × N – 1** zeroes between the upper and lower halves where **K** is the interpolation ratio. The result is K × N spectral values.
8. Perform the Inverse FFT. The result is K × N time history values that are distorted by the Blackman window shape.
9. Select K × n values from the center of the calculated data array.
10. Multiply the result by the inverse of the Blackman Window in the interpolation range.

A REAL DATA SET

Figure 12 shows the results of interpolating a pyrotechnic-shock time history. The data set, that was originally acquired at 1 million samples per second, was modified in the following ways to emulate acquisition with a "low-sample-ratio" converter.

* The data was low-passed at 14 kHz. with a very sharp filter that emulates sigma-delta converter behavior.
* It was then decimated by 32 to produce a 31,250 sample-per-second data set.

This produced a data set with a minimum sampling ratio of about 2.2.

The figure shows the full time history, the fourteen points around the maximum response that were selected for interpolation, and a superposition of the input data, the result, the ideal (undecimated original) and the instantaneous error. The error is less than .01% of peak value.

Smith, "The Interpolation of Sparse Time History Data," Page 6 of 7.

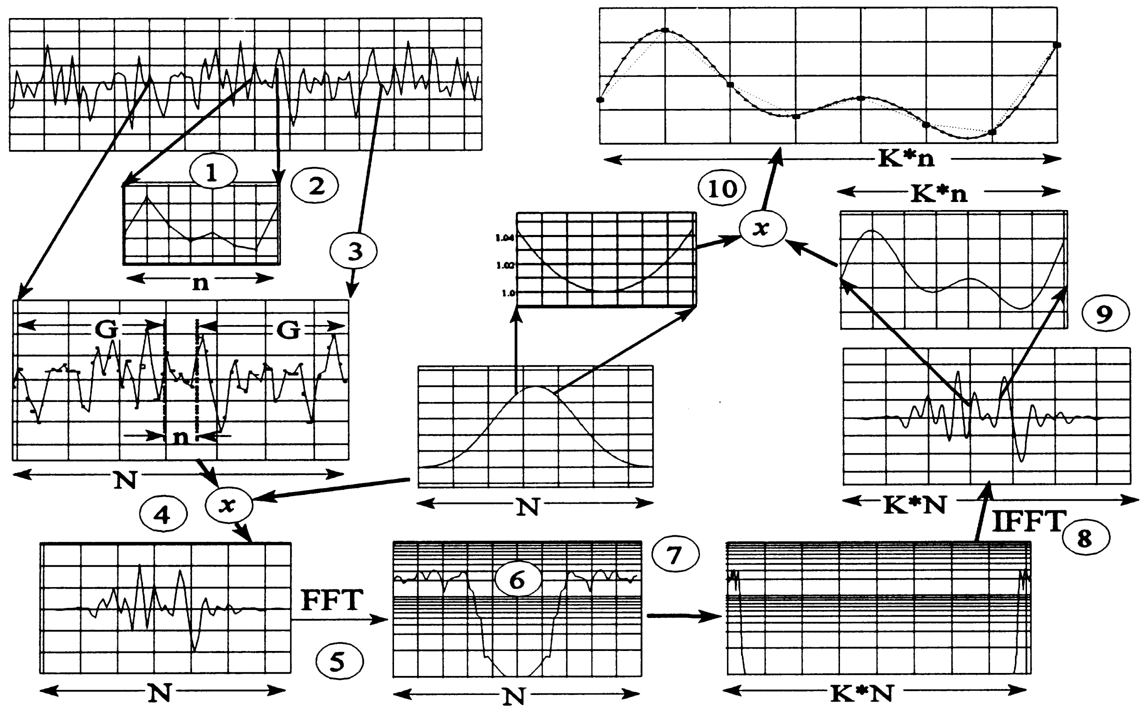

Figure 11 The Interpolation Process

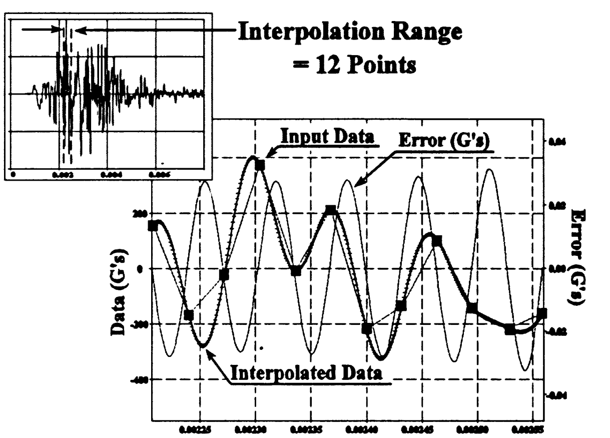

Figure 12. Interpolation (by 32) of a Pyroshock Test Time History.

COMMENTS

There are several features of the interpolation process that have not been directly addressed in the discussion above.

* The procedure described requires that the input signal be band limited (low-pass filtered) so that there are no significant signal components with frequencies above the sample rate/2.1 (which would cause aliasing at a sample ratio of 2.2). Violation will result in significant errors because the window characteristic will not reject the spectrum spreading adequately and, at best, a larger guard band will be needed.

* The interpolation error is not dependent on the interpolation ratio. The worst error occurs half way between the acquired points (interpolation ratio = 2). Higher interpolation ratios only require more calculation.

* The required guard-band size is only loosely related to the interpolation ratio and interpolation range. It is dependent on the sampling ratio but the penalty of handling low ratios is small (guard band of 30 rather than 10 points for ratios of 2.2 vs. 4). It is recommended that the more conservative value be used as the general case.

* An interpolation ratio of 16 is adequate for most applications. With a sample ratio of 2.2, this provides 35.2 points per cycle of the highest frequency component. This, in turn, gives a "worst-peak-determination error" of 0.4% of value. If more accuracy is required, a higher interpolation ratio can be used. The only penalty is computing time.

Technology Training, Inc
E-mail: Training@TTi4edu.com • http://www.TTi4edu.com

Smith, "The Interpolation of Sparse Time History Data," Page 7 of 7.

* The Three-term Blackman window appears to be a good choice for allowed error levels of 0.1%. It offers a good compromise between attenuation at the ends of the time window (which effect the interpolation of larger input ranges) and attenuation of spectral components away from the center frequency. If higher accuracies are necessary, a window with better side-lobe attenuation [6] must be used in conjunction with larger guard-bands.
* The effect of noise on the calculation has not been addressed. It may be argued that, if more points per cycle are acquired, that the data can be analytically smoothed to reduce noise. On the other hand, if lower sample rates are allowed, data acquisition systems with more accuracy and resolution can be used and the data can be smoothed before acquisition to produce the same effect with reduced throughput and storage requirements.
* The results were verified/sanity-checked by experiments using random-input time histories with a variety of interpolation parameters. In all cases, the errors incurred were smaller than those produced by the "cosine" test signal.

CONCLUSIONS

This study demonstrates that the data values between sparsely-acquired points can be calculated to an accuracy of 0.1% or better for sample ratios as low as 2.2.

The procedure described is straightforward and can be implemented on a variety of available data analysis processors. No "special" tools are required.

The capabilities of modern data acquisition systems, that can acquire data with aliasing errors of less than 0.1% when sampling ratios as low as 2.2 are used, can be used to full advantage.

REFERENCES

1. Himelblau, Harry, Piersol, Allan G., et al. 1994. "Handbook for Dynamic Data Acquisition and Analysis," *IES Recommended Practices 012.1*.
2. Smith, Strether, and Hollowell, Bill. 1991. "Techniques for the Normalization of Shock Data," *62nd Shock and Vibration Symposium Proceedings,* Springfield VA.
3. Smith, Strether and Pang, Chao-Sun. November 1993. "Anti-Aliasing Filters for Digital DA Systems," *Sensors Magazine*, Vol 10, No 11, pp 30-33.
4. Race, Randy. March 1994. "Zero Insertion Interpolation," *DSP Applications*.
5. Lynn, Paul A. 1982. *An Introduction to the Analysis and Processing of Signals*, Howard W. Sams & Co, pp 164-166.
6. Nuttall, Albert H., "Some Windows with Very Good Sidelobe Behavior." February 1981. *IEEE Transactions on Acoustics, Speech, and Signal Processing,* Vol ASSP-29, No.1.