Terrigena Ltd.

Time-Gate Filtering in Geophysical Response Estimation

Author:
John E. E. Kingman

Terrigena Technical White Paper Number 2005FEB01.0104(b)
2005 - FEB - 01

Table of Contents

Time-Gate Filters in Geophysical Response Estimation

Usage, design suggestions, and detailed analyses

Overview

Controlled source electrical geophysical time-domain responses are often provided as "time-gates" (also called windows) that reflect moving averages of raw or measured samples of the response half-period waveform. Often, these time-gates are quasi-logarithmically spaced in time and similarly reflect appropriately increasing averaging widths.

By far the most commonly used time-gating scheme entails a uniformly weighted average of all samples in each time-gate. However, there is merit in considering schemes whereby all data are not uniformly weighted in each resulting average. In particular, symmetric tapered shapes in time-gate averaging weights can provide useful and enhanced noise rejection as compared with the standard uniform or boxcar scheme. These advantages and other particulars regarding tapered time-gate designs are addressed below.

In the discussion to follow, the terms "window" and "time-gate" will be used interchangeably.

Fundamentals

Effective Width

We may reasonably view time-gates, or time-gate filtering, as having low-pass characteristics. The frequency-domain character of boxcar or uniformly weighted time-gate filters shows that they suffer poor stop-band attenuation character. This is illustrated in Figure 1, which compares the boxcar and Hanning filter shapes. We see that:

a) the boxcar window enjoys a narrower pass-band than the Hanning window when both are the same length, however, …

b) the Hanning window enjoys deeper stop-band character.

When both windows reflect the same effective lengths; i.e. when the Hanning window is twice the length of the boxcar window, then the Hanning will enjoy generally better characteristics. The term *effective width* is used to consider the fact that samples at the ends of tapered windows are usually rather insignificant in their contribution to the averaged result. We define *effective width* as the sum of all window taps divided by the maximum window tap. The effective width of a boxcar window is then simply the number of taps or length of the window. The effective width of a Hanning window is roughly half the window width.

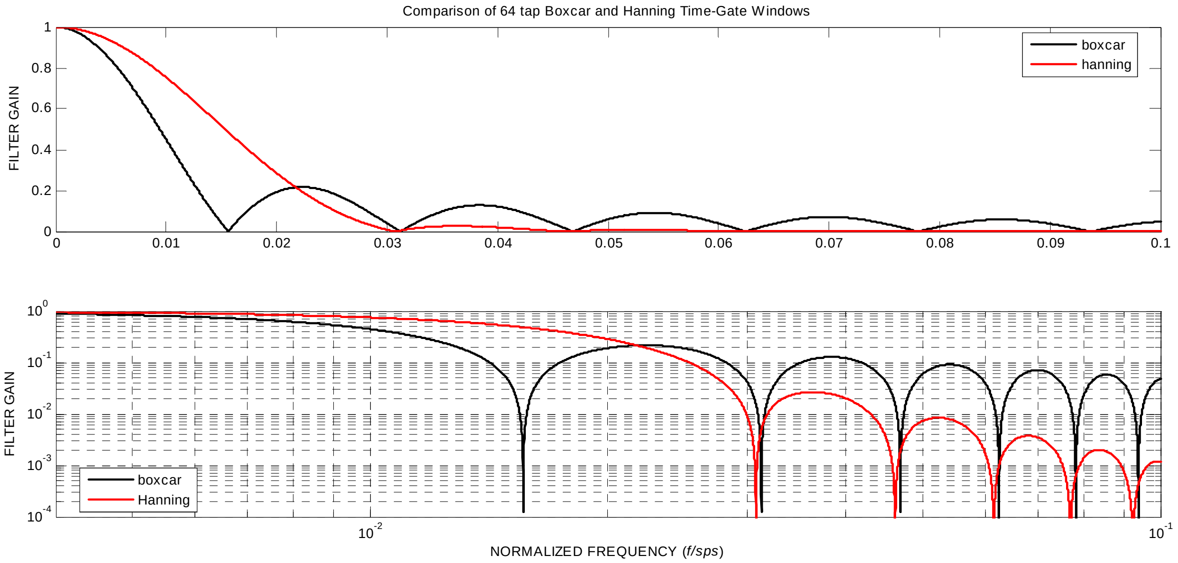

Figure 1 - Frequency domain (amplitude) characteristics of Hanning and Boxcar, or uniformly weighted, filter characteristics. When compared for equal lengths, the boxcar window enjoys a narrower passband (lower cut-off frequency), whereas the Hanning window enjoys a deeper stop-band attenuation floor. For equal "effective widths" (Hanning window x2 wider) the Hanning window enjoys even better advantages.

Figure 2 below illustrates boxcar vs. Hanning character for equivalent effective widths.

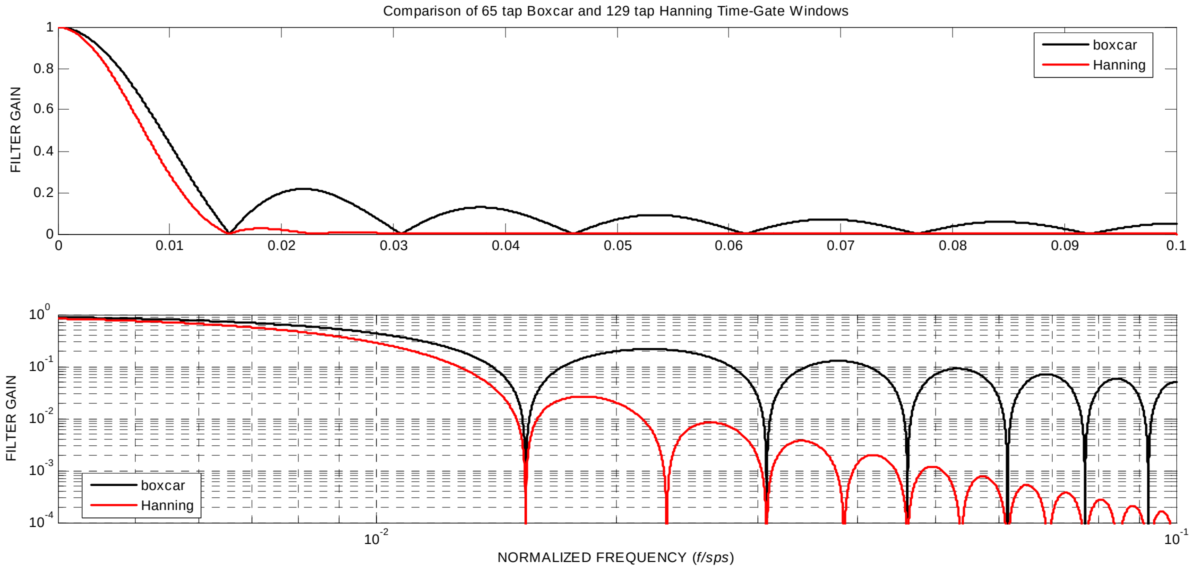

Figure 2 - Frequency domain (amplitude) characteristics for Hanning and boxcar windows with equivalent lengths, in which case the Hanning enjoys a significantly deeper stop-band attenuation floor.

Taper Attribute Details

As suggested, tapered windows enjoy greatly improved stopband characteristics as compared with boxcar windows. By the same token, tapered windows enjoy deeper and/or wider notches in their stop-bands. In general one wishes to align those notches with particularly troublesome noise frequencies such as power-line noise (50 or 60 Hz). Consider a sampling rate of 20,000 sps. 50 Hz noise is then rejected by boxcar windows that are 400, 800, 1200, ... samples wide; i.e. stopband notches will align with 50 Hz and higher harmonics of 50. For Hanning and triangular[1] windows at 20,000 sps, 50 Hz noise is rejected by window widths of 799, 899, 1199, ... samples wide. We illustrate the differences in Figure 3 below:

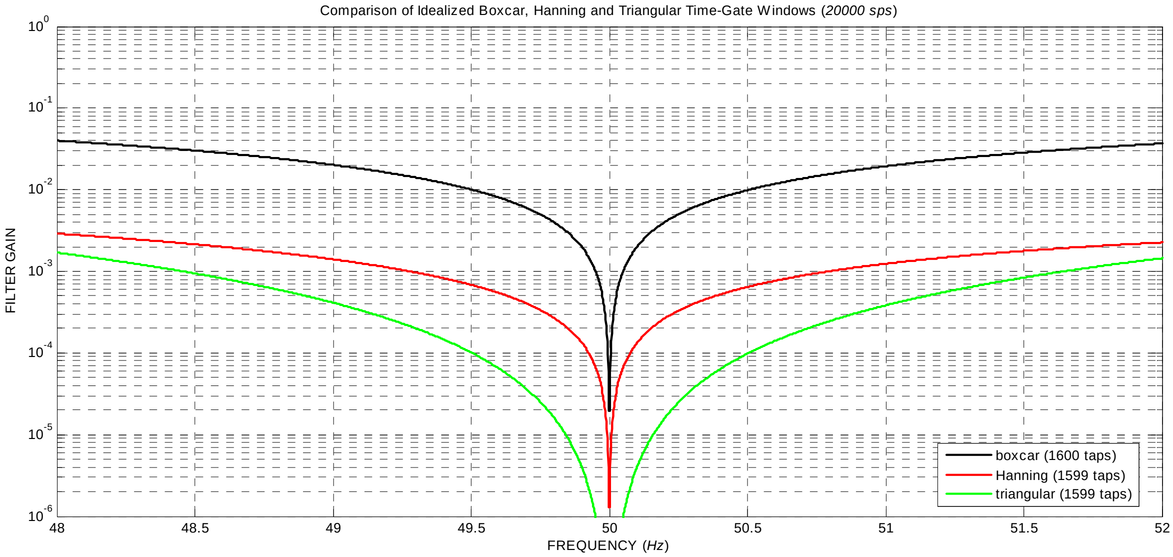

Figure 3 - Frequency domain attenuation notch details, in this case for windows ideally designed to reject 50 Hz noise and higher harmonics. Notice the substantial benefits of the Hanning and triangular window shapes in rejecting noise.

Figure 3 illustrates boxcar, Hanning, and triangular (Bartlett) windows idealized to reject 50 Hz and higher harmonics noise. Of particular interest is the substantially broader and deeper character of the triangular window. The reader should bear in mind that when power-line noise is a problem, careful investigation may frequently reveal a substantial bandwidth to each harmonic of the power-line noise - "substantial" implying perhaps as much as 1 Hz or more. In such cases, when power-line noise is not spectrally narrow to a sufficient degree, then utilizing time-gate designs such as triangular windows that enjoy enhanced widths at the appropriately located notches in the stop-band(s) may prove to be spectacularly beneficial.

While not required, it is generally helpful from a design standpoint if the sampling rate is an integer multiple of the power-line frequency (presuming that such is a dominant noise source). While it is possible to construct idealized windows for rejecting power-line noise when this is not the case, it is considerably less straight forward. Let us presume that we do enjoy a sampling rate equal to an integer multiple of the power-line frequency $\left(sps = n \cdot f_{pwr} \right)$. In this case, idealized window widths are as follows for the indicated windows:

---

[1] Triangular shaped windows are also sometimes referred to "Bartlett" windows.

boxcar: $n, 2n, 3n, \cdots$
Hanning: $2n-1, 3n-1, 4n-1, \cdots$
triangular: $2n-1, 3n-1, 4n-1, \cdots$

The reader should bear in mind that in addition to notches at power-line harmonics, there may be the need to reject other spectrally narrow noise sources such as VLF frequencies, mine site personal emergency device (PED) frequencies, and the like. Late-time windows, where the concern is greatest, may almost always be tailored to reject a given spectrally narrow noise. Two approaches to this are:

1) use a window with a flexible degree of taper and adjust the taper so as to force a stop-band notch at the proper frequency, or

2) use a window that is the convolution two windows, one ideally tailored to reject power-line noise and the other ideally designed to reject the specific noise of interest.

Given that power-line noise is almost always a concern, the latter approach is generally the more powerful. Let us consider a spectrally narrow but troublesome noise at, say, 7413 Hz. For noise between roughly 2/3^{rds} the Nyquist[2] and the Nyquist frequency itself, one of several reasonable approaches to finding an idealized smaller window (that will be convolved with the larger windows) is to try a short (three-tap or slightly more) Hanning window with a baseline offset added to it. In the case at hand, a three-tap Hanning window plus the offset 0.8332 (then normalized for unity gain at DC) provides a strong attenuation notch at 7413 Hz.

$$[0.5, 1.0, 0.5] + 0.8332 = [1.3332, 1.8332, 1.3332]$$

... or properly normalized to unity gain at DC: $[1.3332, 1.8332, 1.3332]/4.4996$

This three-tap window may then be convolved with any other window to reject 7413 Hz noise. When added to (via convolution) the stop-band attenuation of larger time-gate windows this provides a substantial noise rejection capability. This is illustrated in Figure 4 below. Note that convolving the longer Hanning tapered window with the shorter "tuned" window substantially dropped the floor of the stop-band. In effect the Hanning window was tapered more by the convolution. This behavior tracks in general - the greater the taper the greater the stop-band attenuation for equivalent effective widths.

One disadvantage of going to tapers substantially more severe than the Hanning window is that time-gates must become rather wide for a given effective width of interest. This then disallows effective use of the data at the end of a half-period or quarter-period, also moving the latest possible time-gate time (seconds) further back (earlier) in the time-gated response estimates. In general there should be little advantage to using more severely tapered windows than a Hanning window possibly convolved with a 3 to 7 tap shorter window (see discussion regarding standard low-pass FIR designs below). A relatively complete arsenal of useful time-gate window shapes is:

Hanning
Tukey
triangular
boxcar
any of the above convolved with another (usually shorter) window of tuned design

---

[2] The Nyquist frequency is defined as one-half the sampling rate; i.e. sps/2.

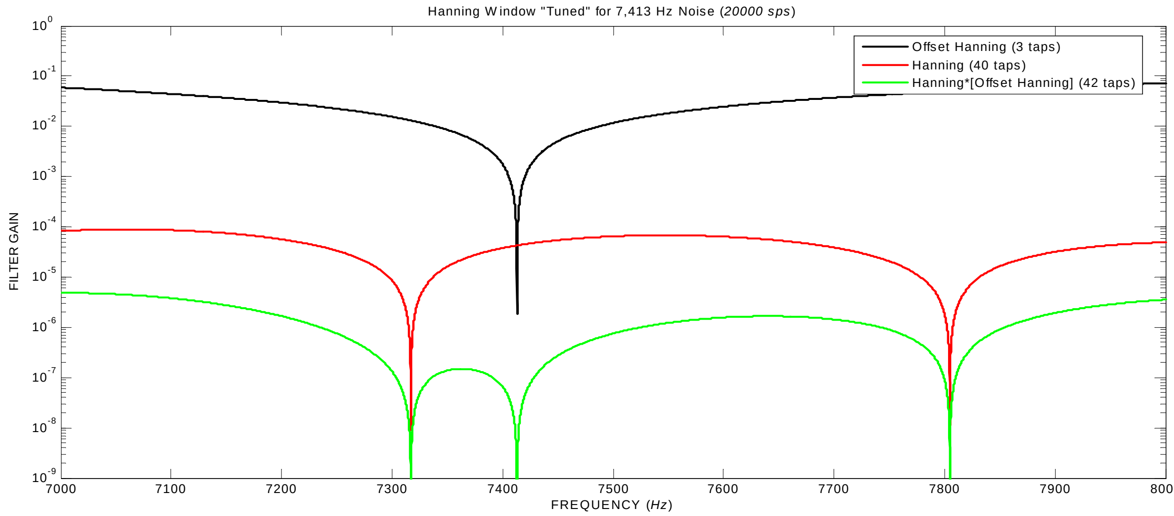

Figure 4 - illustration of a tailored short filter, in this case the three-tap modified Hanning window described above ([1.3332, 1.8332, 1.3332]/4.4996) designed to reject specific frequencies, in this case 7,413 Hz. When the tailored filter is then convolved with standard longer windows a powerful noise rejection character results.

The Tukey window is also sometimes called a "flat-top" window. It is basically a boxcar window with half-cosine tapers at the ends. The cosine tapers may be adjusted in length/width. The Tukey window allows a useful mix of the advantages of the narrower pass-band (for a given window width) of the boxcar window and the deeper stop-band of tapered windows like the Hanning.

As hinted, the Tukey window is adjustable in its shape via a parameter called the *taper ratio* and basically ranges between the boxcar (taper ratio = 0.0) to the Hanning (taper ratio = 1.0). As indicated, the taper ratio dictates the widths of the cosine tapers at either end of the window. The Hanning, triangular and Tukey windows are illustrated below. All have been normalized to ensure unity gain at DC.

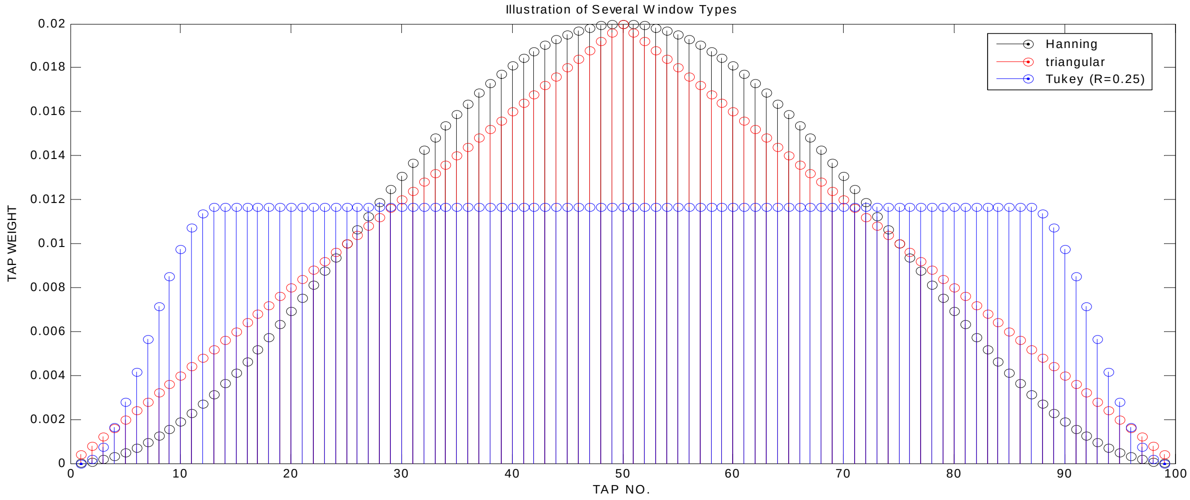

Figure 5 - Examples of 3 useful windows for time-gate filter designs.

The frequency domain characters of variations of the Tukey window reflecting different taper ratios (R) are shown in Figure 6 below.

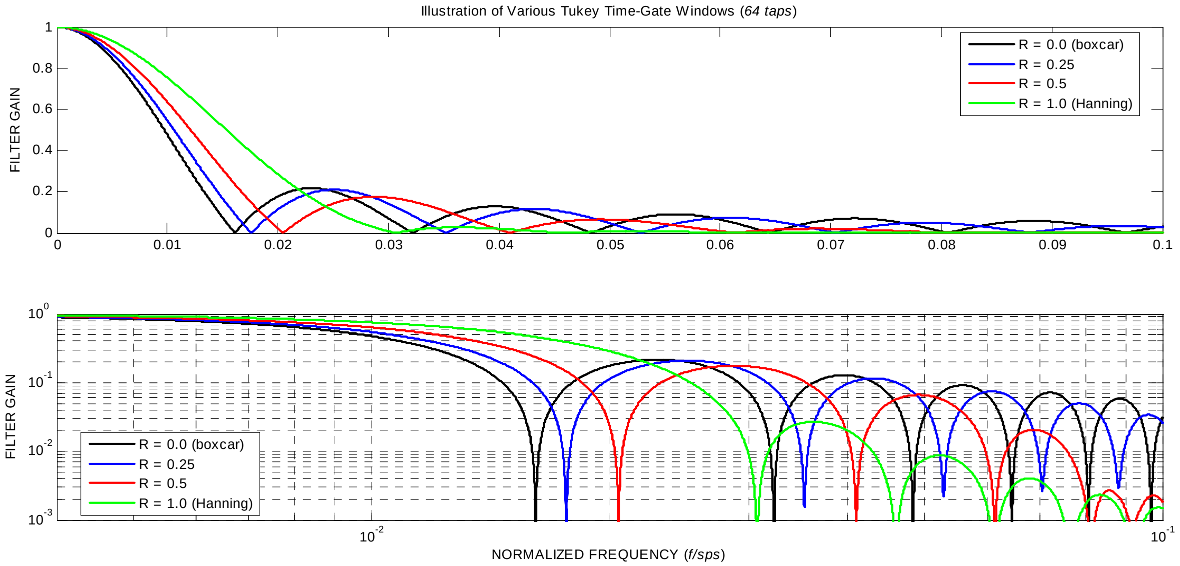

Figure 6 - Illustration of the "continuum" character of the Tukey window ranging from boxcar (no taper) to Hanning. This allows flexibility and optimal tuning to place rejection notches at a given frequency.

Some consideration might be given to time-gate windows that mimic proper low-pass filters as shown in Figure 7 below. Such designs would, in general, provide substantially sharpened low-pass cut-off "walls" and deeper stop-bands, but at the expense of requiring considerably wider windows for a given cut-off frequency as compared with the boxcar window. The disadvantages of such windows in how they impact the latest useable time-gate time were previously discussed. It would seem unexpected if such designs provided substantial benefits as compared with the arsenal of four window designs mentioned above, but the notion should not be discounted off-hand.

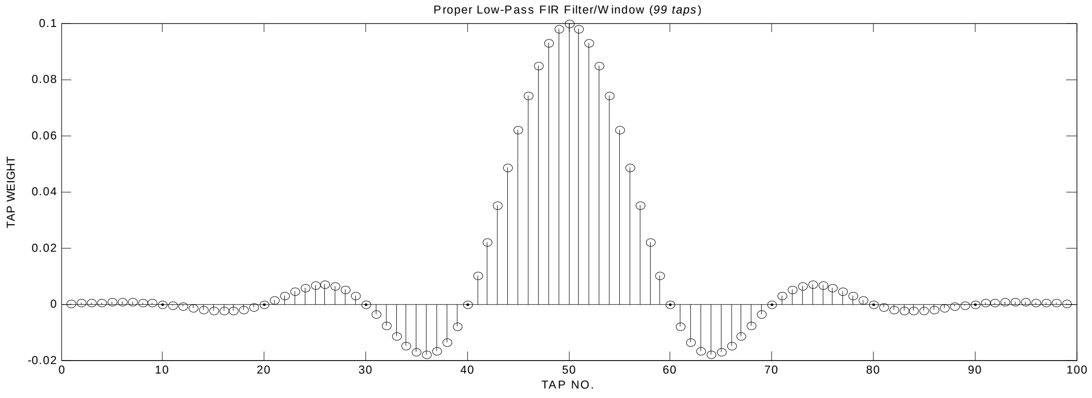

Figure 7 - A "proper" or standard low-pass moving-average (FIR) filter. Like the Hanning as compared with the boxcar window, these require even greater length for a given pass-band width, but essentially provide maximally deep stop-bands if properly designed.

Quasi-logarithmic Expansion Schemes

One rather neat scheme for establishing quasi-logarithmically expanding windows follows a binary design. Consider the following basic binary sequence of window widths: 1, 2, 4, 8, 16, ...

Such a scheme provides rather coarse time resolution, which may be increased by doubling the number of windows with a given width as follows:

1, 1, 2, 2, 4, 4, 8, 8, 16, 16, ⋯

... or for even greater time resolution:

1, 1, 1, 1, 2, 2, 2, 2, 4, 4, 4, 4, 8, 8, 8, 8, 16, 16, 16, 16, ⋯

In this scheme we use binary window widths and a resolution factor that designates the number of repeats for each window size. Doubling the resolution factor tends to (roughly) double the time-gate time resolution or number of time-gates in a response quarter-period or half-period.

Odd-valued window lengths are often needed or preferred, which may be provided using a simple adjustment to the binary scheme: either add or subtract one sample (discounting any resulting zero-lengths), as follows using a resolution factor of 4:

1, 1, 1, 1, 3, 3, 3, 3, 5, 5, 5, 5, 9, 9, 9, 9, 17, 17, 17, 17, ⋯

1, 1, 1, 1, 3, 3, 3, 3, 7, 7, 7, 7, 15, 15, 15, 15, 31, 31, 31, 31, ⋯

A suggested generalized approach for designing time-gate expansion patterns is:

1) For the taper or targeted window type and sampling rate, determine the precise window lengths that perfectly reject power-line noise.

2) Incorporate any adjustments or tuning required via convolution with a shorter filter/window (such as explained in the exercise illustrated in Figure 4 above where a filter for 7,413 Hz was designed).

3) Incorporating the ideal window lengths as determined by steps 1) and 2), meld those windows into a quasi-logarithmic binary scheme as described above. For example, let's presume that the ideal window lengths for power-line noise considerations are 799, 899, 1199, 1599, etc. Let's further presume that we want to "tune" to add a particular notch in the stop-band using a 5-tap window. Convolving the 5-tap window with the longer ideal windows yields lengths of 803, 903, 1203, 1603, etc. taps. We generally would start with the shortest ideal window and expand with the binary scheme as far as deemed needed. If a resolution factor of 4 is chosen, then we'd design the time-gate scheme to continue with windows of the following lengths - carried as far out as needed: 803, 803, 803, 803, 1603, 1603, 1603, 1603, 3203, 3203, 3203, 3203, ...

4) Next we fill in the earlier times by extrapolating the same sequence backwards until single-value windows are encountered (in reverse order). To the degree possible these would continue to use the "tuning" adjustment (via convolution) of step 2). Example:

    403, 403, 403, 403, 203, 203, 203, 203, 103, 103, 103, 103, 53, 53, 53, 53, …

27, 27, 27, 27, 13, 13, 13, 13, 7, 7, 7, 7, 3, 3, 3, 3, 1, 1, 1, 1

5) Pad the early times with as many single-valued windows as deemed appropriate

6) Next we need incorporate any needed overlap between windows. Ideally one desires a uniform weighting of all data/samples in the response. Odd-valued Hanning, and triangular windows of equal lengths may generally be overlapped by roughly half their length to ensure perfectly uniform data weighting. Similarly, the Tukey window usually may be overlapped by half the taper width to yield uniform weighting. Obviously, however, when adjacent tapered time-gates are of different lengths then it is impossible to achieve uniform data weighting presuming the windows are symmetric. In this regard a reasonable compromise is the best that can be achieved.

7) Lastly, using the window overlap scheme determined in step 6) we calculate the sample index ranges for all windows up to the last window that does not exceed the quarter-period or half-period boundary.

As a general rule, the coarser the time-resolution (or resolution factor in scheme under discussion), the larger the ending time-gate(s) will be. Clearly at some point the time-gates will provide a significantly distorted representation of the signal if too large. So the question remains, how large is too large? Figure 8 and Figure 9 below show time-gated results for theoretical step-function (turn off) in-loops EM responses over a homogeneous half-space. The time-gate scheme used straight binary with a resolution factor of 1.0; i.e. with window widths of: 1, 2, 4, 8, 16, ... samples, and overlapping roughly 1/2 window width. A sampling rate of 100,000 sps is presumed but no anti-aliasing filter effects are incorporated - infinite bandwidth of the sampled signal is presumed.

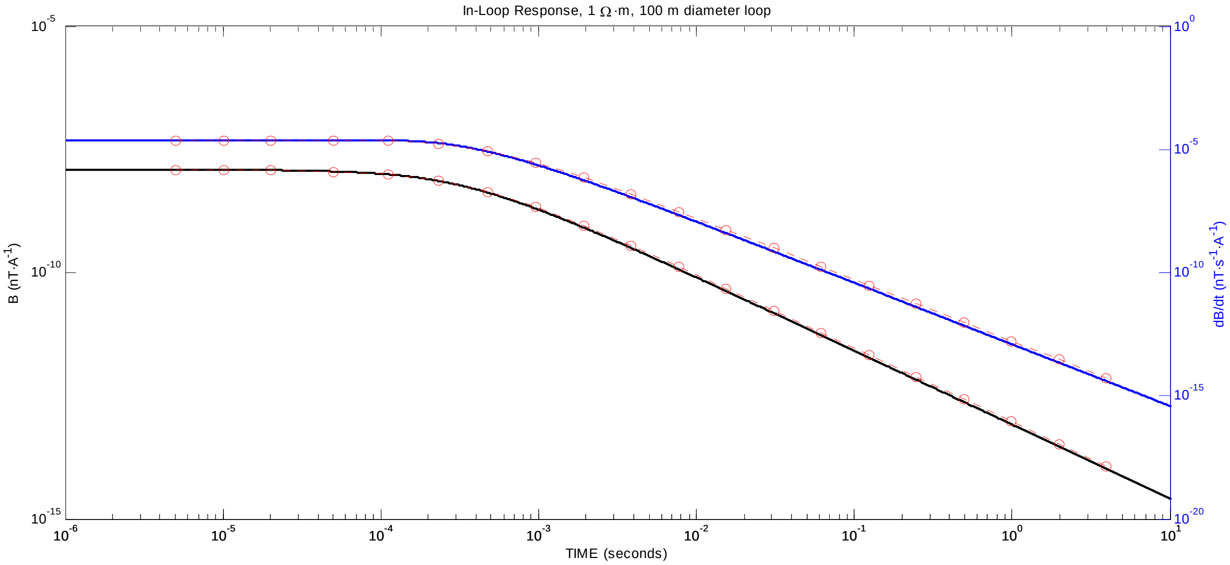

Figure 8 - Time-gate filtering of an In-Loop EM response using coarsest time-resolution under the binary scheme, with window widths increasing with each time-gate (resolution factor = 1). There is a slight offset (on a log ordinate scale) suggesting that the time-gates are a bit too wide for their respective time positions. This problem is best relieved by increasing the resolution factor. The sampling rate modeled is 100,000 sps but distortion from band-limiting processes such as anti-aliasing filtering is not incorporated; the implied bandwidth is infinite.

Figure 9 repeats the experiment of Figure 8 but for a more resistive, 100 Ω•m earth. In this case it is seen that few additional single-sample time-gates are needed.

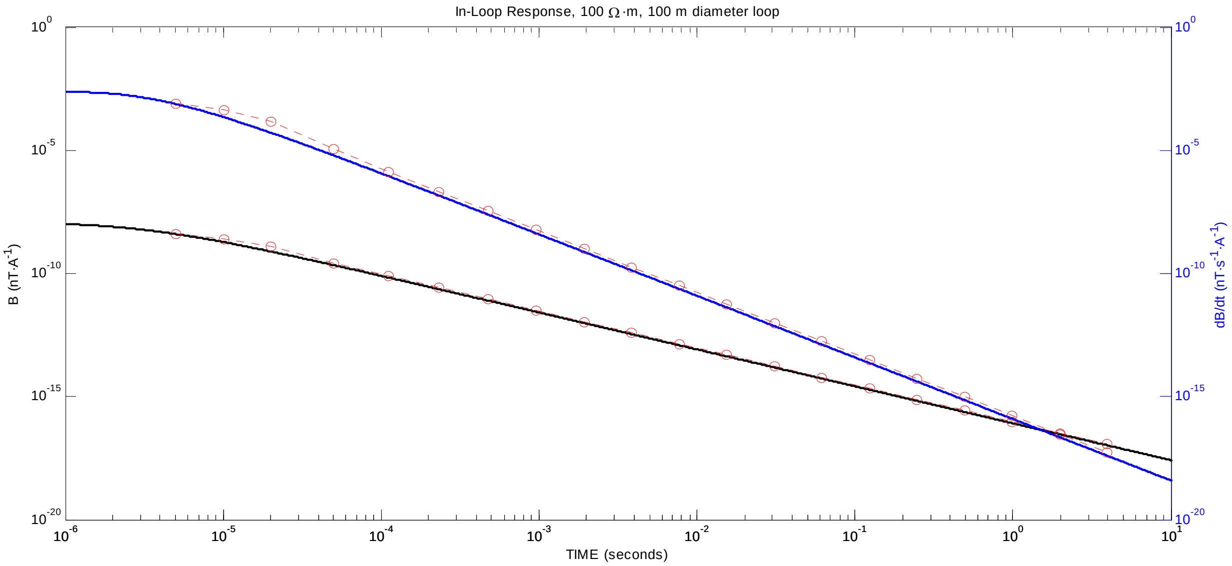

Figure 9 Time-gate filtering of a theoretical step response, just as in Figure 8 (resolution factor = 1) but wherein the ground is more resistive. Here again we see evidence of distortion related to the excessive time-gate widths. The bump in the early times, most obvious for time-gate no. 3, is readily fixed by padding the gates with additional single-sample time-gates. The sampling rate modeled is 100,000 sps but distortion from band-limiting processes such as anti-aliasing filtering is not incorporated; the implied bandwidth is infinite.

In Figure 8 and Figure 9 above, clearly there is a noticeable (at this visual scale) degree of distortion owing to the time-gate widths as compared with the curvature of the response. Carrying the exercise further with greater time resolution and additional earth resistivities (results not illustrated) we find that:

- In general additional single-sample padding is prudent - the degree of such being dependent on the details of the actual transmitted waveform, response geometry, etc. Also, in real cases with finite switching ramps a critical consideration is when t = 0 is defined; i.e. whether at the start or end of a current switching ramp.

- In the binary expansion scheme a resolution factor of 1 (i.e. non-repeated window widths of 1, 2, 4, 8, 16, ..., etc.) is a bit coarse. Resolution factors of 2 or more are strongly recommended.

For white noise we expect the improvement in SNR to track the square-root of the window effective width. In general for realistic conditions noise may be expected to be spectrally colored, and the improvement in SNR may often not be as good as expected for white noise. Regardless, if late-time data are suffering poor repeatability, there may be good reason to widen the later time-gates to improve SNR. The degree to which such will distort the signal as compared against un-filtered responses will depend much on factors mentioned earlier as well as ground conditions. **The safe and appropriate method of dealing with the issue of time-gate filtering related distortion is to impose the same time-gate process used in estimating responses for forward modeling and/or inversion calculations.**

A reasonable rule-of-thumb to help time-gated EM responses track theoretical infinite bandwidth responses adequately is to ensure that the maximum time-gate effective width is 1/8th a quarter-period or less. There are circumstances that could make this rule inadequate. Again, the appropriate and robust way to deal with the issue is to ensure that the theoretical calculations of forward or inverted models use the same time-gate scheme and bandwidth as apply to the measured data and response estimates.

Aliasing

Time-gate filtering does not circumvent the problems of aliasing. As suggested by Figure 10, noise at nearly integer multiples of the sampling rate is not attenuated by any moving average (FIR) filtering process. These frequencies fold back to roughly DC - exactly DC at exactly an integer multiple of the sampling rate. Hence time-gate filtering will not prevent aliased noise sources from appearing as low frequency noise in the response estimates. The electrical geophysical community seems to have taken a cavalier and risky approach to the issue of aliasing in measured data. This is ill-advised and, in the author's opinion, based on presumptions that cannot possibly be assured in all cases. Good geophysical receivers will always ensure against aliasing in the analog-to-digital conversion process. As mentioned earlier, the impact of limiting the bandwidth of the measured signal will generally need to be modeled or accounted for in the interpretation process.

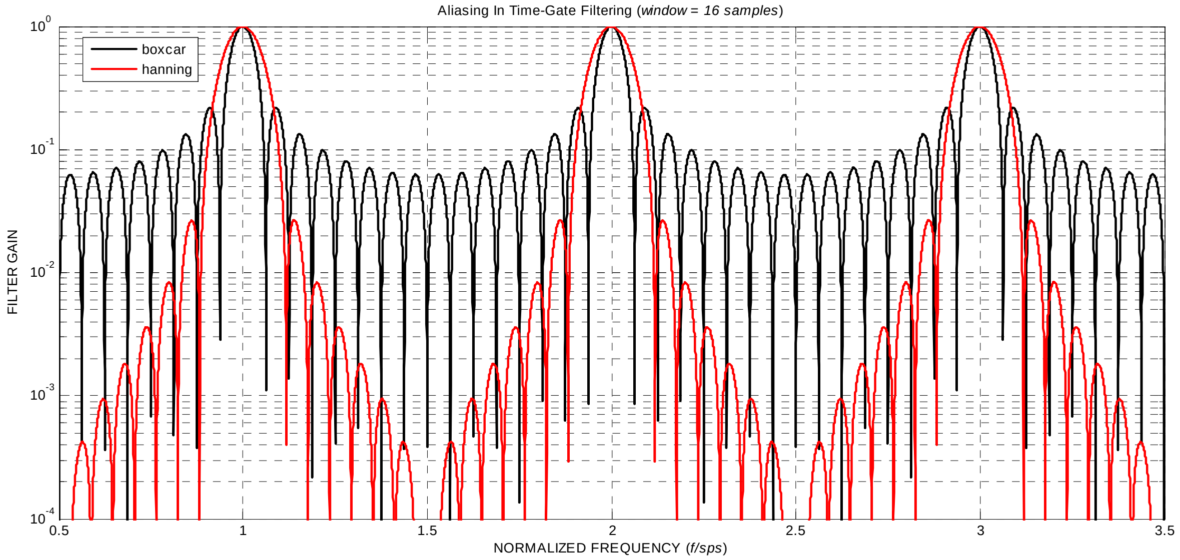

Figure 10 - Time-gate window frequency domain character for frequencies greater than the sampling rate, showing that time-gate filtering does not relieve problems associated with aliasing; i.e. time-gate windows readily pass noise at frequencies corresponding to multiples of the sampling rate. If such noise sources existed they would wrap back so as to be seen as low frequency noise. Receiver design practices that do not ensure against aliasing are ill advised.

References

Mathworks, 2004, *Signal Processing Toolbox - User's Guide, Version 6*, The Mathworks

Hamming, R.W., 1977, *Digital Filters*, Prentice-Hall, Inc.

Author Contacts

John E. E. Kingman
john.kingman@terrigena.com