Understanding Invisible Clipping and the Need for Distributed Gain

Thomas P. Gerber, Douglas R. Firth, Alan R. Szary

Precision Filters, Inc.

Ithaca, New York
(607) 277-3550
E-mail: sales@pfinc.com

Key Takeways

- If an amplifier in a signal conditioner is overloaded and saturates on its input, the amplifier output will be clipped and the data corrupted.
- If this clipped output is then passed through a low-pass filter, evidence of the clipping is removed – it becomes invisible – but the in-band signal will be a distorted version of the true signal.
- The risk of invisible clipping can be minimized by distributing gain across both sides of the filter, a standard feature on Precision Filters' signal conditioning products.
- By providing software-programmable apportionment of gain, Precision Filters makes it easy for users to prevent invisible clipping from corrupting in-band signals.

Introduction

Analog signal conditioners amplify low-level input signals to ensure adequate dynamic range for digital conversion. Because the components of the signal conditioning circuit add noise that accumulates along the signal chain, it is advisable to apply gain early in the signal flow to maximize signal-to-noise ratio. For this reason, most commercially available amplifiers place all of the amplifier gain before the low-pass filter. Often the gain is simply set to amplify the expected sensor full-scale output to match the input level of the recording device. However, if the signal is contaminated with out-of-band noise from ambient sources or resonance of the sensor itself, then applying gain to the total signal (in-band plus out-of-band components) can cause the amplifier to saturate. Saturation will clip the signal to the maximum output level of the amplifier. Subsequent passage though the low-pass filter will remove evidence of the clipping – the flat-topping is no longer visible – but the waveform at the output of the filter will be a distorted version of the true in-band signal. The clipping is therefore invisible, since there is no way to know if clipping has occurred from visual inspection of the output waveform.

This paper demonstrates the effects of invisible clipping using a simple test signal. Strategies for preventing invisible clipping are then discussed, including the use of distributed gain, programmable gain reserve, and overload detection – features that are standard in all Precision Filters (PFI) signal conditioning products.

Invisible Clipping: A Test Case

To illustrate the effects of invisible clipping on signal conditioner output, consider a stationary test voltage signal from a hypothetical sensor with the following form:

$$y_s = y_{ib} + y_{ob} \tag{1}$$

where y_s is the total signal, y_{ib} is the in-band signal of interest, and y_{ob} is out-of-band noise. The in-band signal is modeled using the first two (odd) harmonics of a square wave,

$$y_{ib} = A\sin\left(2\pi f_1 t\right) + \frac{A}{3}\sin\left(2\pi f_3 t\right),\ f_3 = 3f_1 \tag{2}$$

The frequencies f_1 and f_3 are set to 5 kHz and 15 kHz, respectively, and the amplitude A is set equal to 1.061 V to yield a peak in-band signal level of 1 V (Figure 1). This signal is contaminated with a source of out-of-band white noise (y_{ob}) that is band-limited with a center frequency of 100 kHz (Figure 1B) and an rms voltage of 2.4 V. It is assumed that the presence and level of the out-of-band noise y_{ob} is unknown at the time of measurement.

Now assume the signal y_s in Figure 1 is fed into a signal conditioner and amplified prior to low-pass filtering. The desired peak signal level for input to the downstream analog to digital converter (ADC) is 4 V. The gain on the amplifier is therefore set to x4 to yield an in-band signal with a peak level of 4 V. However, the gain is also applied to the out-of-band noise level, resulting in a total amplified signal that exceeds the ±10 V limit on the amplifier output (Figure 2). Consequently, the total signal y_s saturates the amplifier and the signal is clipped.

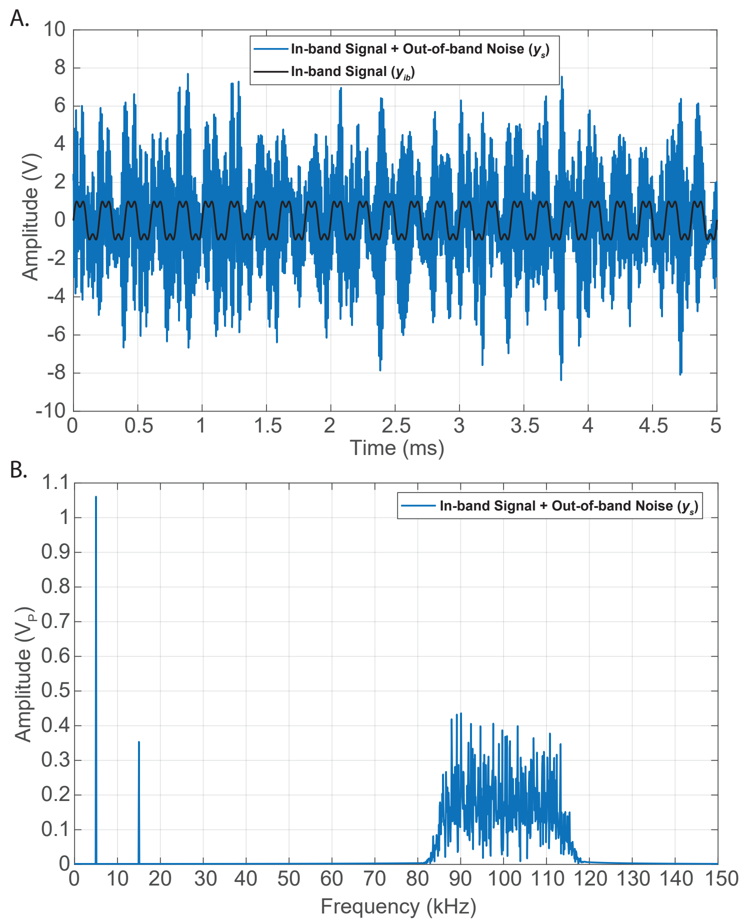

Figure 1. (A) Test signal constructed to demonstrate the effect of invisible clipping. The total signal y_s includes an in-band component y_{ib} (peak voltage = 1 V) and out-of-band noise y_{ob} (rms voltage = 2.4 V). (B) Amplitude spectrum of y_s. The in-band square wave components f_1 (5 kHz) and f_2 (15 kHz) are clearly distinguishable from the band-limited noise centered at 100 kHz.

To simulate passage of this clipped signal through a subsequent low-pass filter stage, PFI's 8-pole LP8P filter transfer function is applied to it. The LP8P is similar to a Bessel filter: they both have a linear phase response, but the amplitude roll-off of the LP8P is sharper than an 8-pole Bessel. The LP8P is set to a cutoff frequency (F_C) of 25 kHz to ensure preservation of the in-band signal y_{ib} and rejection of the out-of-band noise y_{ob}.

Phone: 607-277-3550
E-mail: sales@pfinc.com

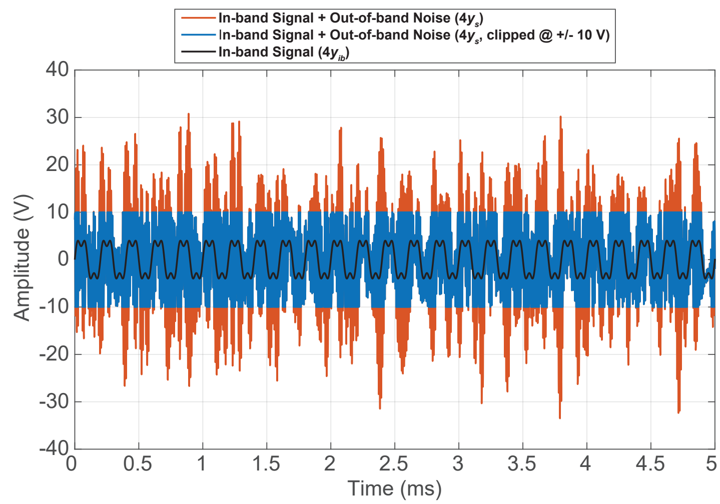

Figure 2. Test signal after passage through the pre-filter amplifier (gain = x4). Beyond the ±10 V clipping threshold, the total amplified signal (= 4 y_s) is clipped and limited to a peak voltage of 10 V. The amplified in-band portion of the signal (= 4 y_{ib}), which lies below the clipping limits of the amplifier, is shown for reference.

The resulting filter output is shown in Figure 3A. The model in-band signal y_{ib} is also shown with a gain of x4 applied for comparison. Notice that the filtered waveform is a distorted version of the amplified in-band signal, despite the clear spectral separation between the in-band signal and out-of-band noise. Moreover, there is nothing in the waveform to indicate the input to the filter was clipped. Importantly, this distortion of the in-band signal occurs even though the clipping threshold of the amplifier (± 10 V) is above the maximum in-band signal level (4 V). The distorting effect of clipping on the in-band waveform can be understood by comparing amplitude spectra between clipped and unclipped filter inputs (Figure 4). Clipping introduces spurious spectral energy into the filter's passband (Figure 4B) and reduces the magnitude of the in-band signal components (i.e. f_1 and f_3).

For comparison, Figure 3B shows the output of the filter when the input signal is not clipped. In this case, the LP8P is applied directly to the output from the sensor (i.e. y_s, as shown in Figure 1), with the gain applied to the filter output. The resulting waveform matches the in-band signal y_{ib}, with slight attenuation and time delay that are determined by the LP8P's frequency response.

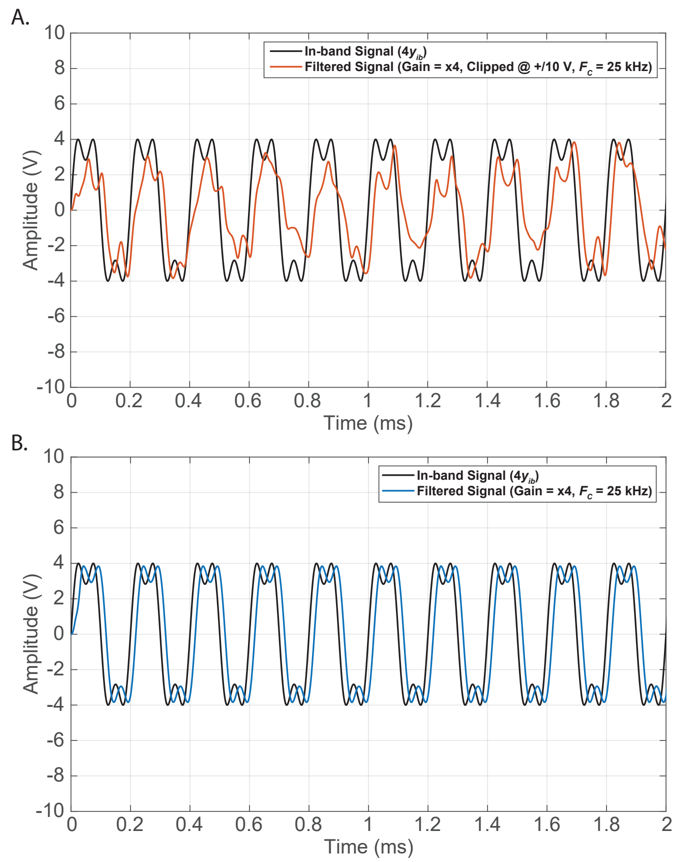

Figure 3. (A) Output (0-2 ms) produced by filtering the amplifier-clipped total signal (i.e. $4y_s$ clipped at ± 10 V) with PFI's LP8P filter set to a cutoff frequency (F_C) of 25 kHz. The amplified in-band signal (i.e. $4y_{ib}$) is shown for comparison. (B) Output (0-2 ms) produced by filtering y_s with PFI's LP8P filter set to a cutoff frequency (F_C) of 25 kHz. In this case all gain (x4) is applied after the filter. The amplified in-band signal (i.e. $4y_{ib}$) is shown for comparison.

Distributed Gain and Gain Reserve

The preceding example demonstrated how clipping by a pre-filter amplifier can degrade the in-band signal in the filter output. The clipping occurred because the input to the pre-filter amplifier contained out-of-band noise that was either unknown or poorly constrained. This degradation was shown not to occur if the amplification was reserved for a post-filter stage. It is tempting to conclude from this that the gain required for the in-band signal should always be applied after any filter stages.

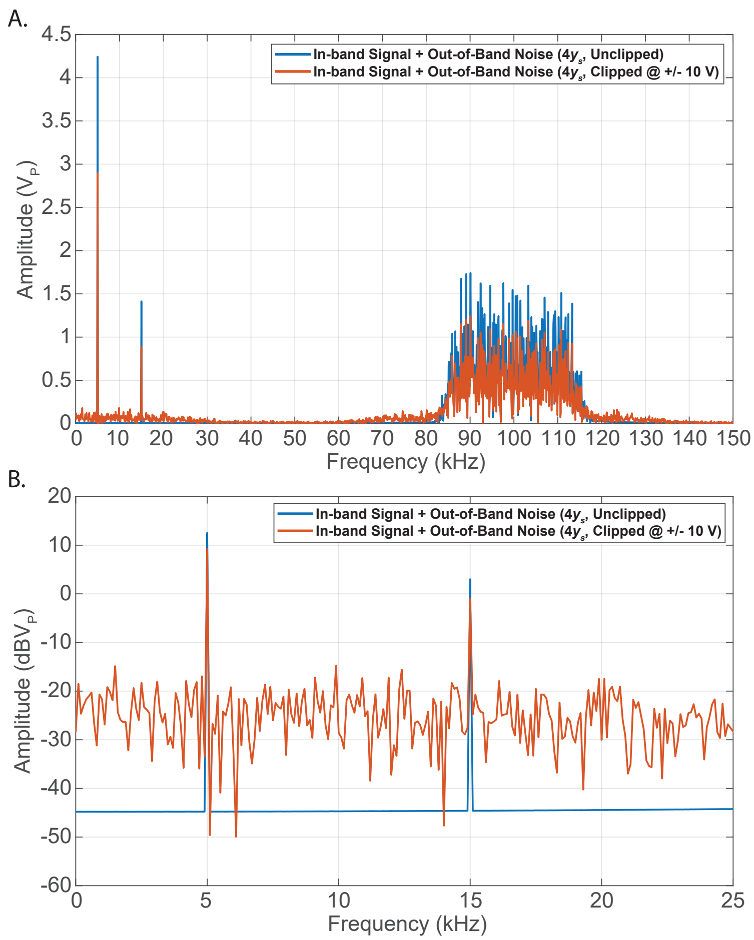

Figure 4. (A) Comparison of amplitude spectra between clipped and unclipped versions of the amplified signal y_s. (B) Inset from (A) showing the passband of the filter used in this study (LP8P, $F_c =$ 25 kHz) on the decibel scale.

But this approach ignores the fact that the op-amps and filters comprising the signal conditioning chain introduce noise as well. Though generally low-level, if this self-noise is picked up and subsequently amplified by a later gain stage, it can significantly reduce the signal-to-noise ratio (SNR) of the recorded signal. For this reason, it is advisable to apply as much gain as possible close to the sensor.

With this tradeoff in mind, Precision Filters offers signal conditioning topologies that distribute the gain needed for the in-band signal across both sides of the low-pass filter stage. In this way, a balance can be struck between preserving SNR and protecting against signal clipping due to out-of-band energy. The pre-filter gain (PRG) and post-filter gain (POG) can be independently programmed to achieve a desired total gain. To facilitate setting the PRG and POG, it is also helpful to define the gain reserve, which can be used to quantify the required headroom for out-of-band noise to avoid clipping by the pre-filter amplifier.

To understand the gain reserve, consider the signal flow depicted in Figure 5. From (1), we can define the input to the low-pass filter, y_{in}, as

$$y_{in} = G_{PR}\left(y_{ib} + y_{ob}\right) \tag{3}$$

where G_{PR} is the pre-filter gain factor (PRG). Now assume the filter removes the out-of-band noise, preserving only the in-band component so that the output from the post-filter amplifier y_{out} is

$$y_{out} = G_{PO}\left(G_{PR}\, y_{ib}\right) = G_T\, y_{ib} \tag{4}$$

where G_{PO} is the post-filter gain factor (POG) and G_T is the total gain factor. Dividing (3) by (4) gives

$$\frac{y_{in}}{y_{out}} = \frac{G_{PR}}{G_T}\,\frac{y_{ib} + y_{ob}}{y_{ib}} = \frac{R}{G_{PO}} \tag{5}$$

where the gain reserve R is defined by

$$R = \frac{y_{ib} + y_{ob}}{y_{ib}} = 1 + \frac{y_{ob}}{y_{ib}} \tag{6}$$

Between (5) and (6), it should be apparent that if the maximum value of y_{in} is determined by the clipping threshold of the pre-filter amplifier, and the maximum value of y_{out} is determined by the input limits on the recording device, then the gain reserve R defines how large the in-band signal plus out-of-band noise can be relative to the in-band signal before clipping will occur. (Or, equivalently, $R-1$ defines how large the out-of-band signal can be relative to the in-band signal.) With (5), expressions for PRG and POG can be written in terms of R:

$$G_{PR} = \frac{1}{R}\,\frac{y_{in}}{y_{ib}} \tag{7}$$

$$G_{PO} = R\,\frac{y_{out}}{y_{in}} \tag{8}$$

For example, consider an in-band signal from a sensor with a peak voltage of 5 mV. This signal will be digitized by an ADC with a maximum input of 5 V. To maximize dynamic range, a total gain G_T of x1000 is desired for the in-band signal. To guard against out-of-band noise, a gain reserve of 4 is chosen, allowing for out-of-band noise to exceed the level of the in-band signal by a factor of 3 without clipping. The clipping threshold of the pre-filter amplifier (± 10 V) sets the maximum value of y_{in} at 10 V. From (7), this equates to a

Phone: 607-277-3550
E-mail: sales@pfinc.com

PRG of x500 and thus a POG of x2. With these settings, the in-band signal will be amplified to a peak level of 2.5 V prior to filtering while the out-of-band signal will be amplified to a peak level of 7.5 V, ensuring the clipping threshold of 10 V is not exceeded.

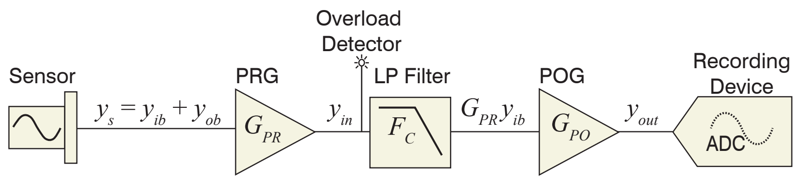

Figure 5. Signal flow diagram for a signal conditioning system with a distributed gain topology

In practice, a software GUI allows users to control and set the gain reserve employed on PFI's signal conditioners, with options to independently set PRG, POG, gain reserve, and the maximum levels of y_{ib} and y_{out}. Because the gain resolution of the amplifiers limits the PRG and POG to discrete levels, PFI's software calculates and displays the actual gain reserve for any combination of the above parameters so that the actual out-of-band protection is known at the time of measurement.

Even with the option of programming gain reserve using distributed gain, there may be instances where the anticipated out-of-band noise levels are underestimated. In this case a signal may clip even with out-of-band protection. For this reason, PFI places overload detectors between the low-pass filter and pre-filter amplifier to alert users when an amplifier has saturated. This provides an extra layer of protection, allowing users to verify that their recorded data has not been corrupted by otherwise invisible clipping at a pre-filter stage.

Summary

Invisible clipping can be an insidious problem. Like aliasing, the corrupting effect of clipping cannot be easily identified and removed in a sampled signal. So in addition to applying low-pass anti-alias filters prior to sampling, care must be taken to ensure that amplifier gain is set so that clipping due to out-of-band energy is prevented. Signal conditioning topologies that distribute gain across both sides of the low-pass filter provide a way to simultaneously optimize SNR and minimize the risk of amplifier clipping. By providing programmable PRG, POG, and gain reserve, along with built-in overload detection, PFI's signal conditioners are designed to ensure valid data can be collected even when the in-band signal from a sensor is dominated by out-of-band noise.

Precision Filters, Inc.
Ithaca, New York
(607) 277-3550
E-mail: sales@pfinc.com

PRECISION FILTERS, INC.
Transducer Conditioning Systems ■ Filter/Amplifier Systems ■ Signal Switching Systems