Filter Response of the PARxCH Family of A/D Converters

Symmetric Research Application Note AN0020

July 19, 2003

Introduction

The PAR1CH, PAR4CH and PAR8CH analog-to-digital converter boards from Symmetric Research utilize the Burr-Brown ADS1210 chip, a second-order sigma-delta A/D converter incorporating a three-stage low-pass filter. The digital low-pass filter is an essential component of the ADS1210 (and in general of all sigma-delta converters) and defines the precision of the output digital samples. In this application note, we describe the theoretical filter response of the ADS1210, and experimentally verify this response by evaluating the digital output from known, accurate square-wave and triangle-wave analog input signals.

Sigma-Delta A/D Converters

A generic sigma-delta converter consists of the following components:

1. A *modulator* (a high-frequency single-bit sampler),
2. A low-pass *digital filter*,
3. A *decimator* that subsamples to the desired output rate.

The modulator samples the analog input signal at a rate f_{MOD} typically much higher than the desired output rate f_{OUT}. Each sample from the modulator takes on one of two values: the maximum and minimum voltage levels that the A/D converter may accept without saturating. The modulator thus generates a time series of 1-bit samples, but does so in such a way that the original input signal is recoverable from the low-frequency components of the time series. To understand how this magic can occur, let's examine the modulator stage in more detail.

Figure 1 illustrates the sigma-delta modulator implemented in the ADS1210. The *order* of the modulator is defined by the number of differencing (delta) and integration (sigma) units before the 1-bit quantizer. The quantizer rounds the signal to the nearest voltage level V_{MIN} or V_{MAX}. This nonlinear operation is theoretically modeled by assuming that a quantization signal Q has been added to the input to form the output signal Y, which can take on one of two possible values, V_{MIN} or V_{MAX}.

A Z-transform analysis of the circuit in figure 1 is easily performed. A one-sample delay built into the 1-bit DAC unit is modeled by the operator z, and the integration unit is modeled by

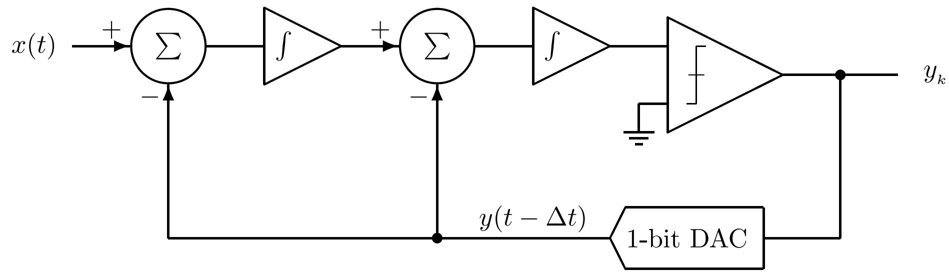

Figure 1: Block diagram of a second-order sigma-delta modulator. The input signal $x(t)$ is fed through two stages, each consisting of a differencer (the *delta* or $\sum$ unit) and an integrator (the *sigma* or $\int$ unit). The comparator examines the resulting signal and assigns a 0 to the sample if it is closer to V_{MIN}, or 1 if it is closer to V_{MAX}. The digital output stream y_k consists of 0's and 1's. The sampling rate of the comparator is typically many hundred times higher than the desired output sampling rate of the A/D converter. The 1-bit DAC unit converts the sample back to V_{MIN} or V_{MAX} (the signal $y(t)$ in the text), imposing a one-sample delay in the process, and then negatively feeds the signal back to the delta units.

$1/(1-z)$. The Z-transform expression for the circuit is then

$$\left[(X(z) - z\,Y(z))\,\frac{1}{1-z} - z\,Y(z)\right]\frac{1}{1-z} + Q(z) = Y(z) \tag{1}$$

Solving for $Y(z)$ yields

$$Y(z) = X(z) + (1-z)^2 Q(z) \tag{2}$$

where $z \equiv \exp(i\omega\Delta t_{\text{MOD}})$ and Δt_{MOD} (the modulator sample interval) $\equiv 1/f_{\text{MOD}}$.

The fundamental assertion in sigma-delta conversion is that the quantization signal Q is a random, independent time series. For certain special input signals, this assumption may fail, for example in the case of idle tones [3], but otherwise it holds remarkably well. If so, the amplitude spectrum of $Q(z)$ is uniform, with amplitudes of the same order as $X(z)$, and the effect of the filter $(1-z)^2$ in equation 2 is to suppress the low frequency components of Q relative to X. That is, for values of z close to unity, $Y(z) \approx X(z)$. To recover X from Y, simply apply a low-pass filter to Y. For more details on the theory of sigma-delta converters, see reference [1].

Low-Pass Filtering and Decimation

The output time series is derived by filtering the modulator time series, then decimating the series by a whole number of samples. The output sample rate f_{OUT} is therefore controlled by the modulator sample rate f_{MOD} and the decimation factor N (a whole number):

$$f_{\text{OUT}} = f_{\text{MOD}}/N. \tag{3}$$

For the ADS1210, N may range from 19 to 8000 [2]. Given the decimation factor N, the simplest way to filter and decimate is to form the average of every group of N input samples.

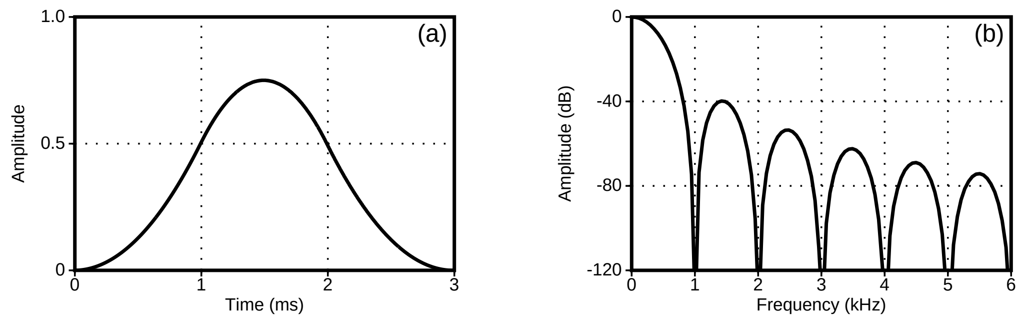

Figure 2: Three-stage low-pass averaging filter f_3 (equation 6). In this example, $f_{\mathrm{MOD}} = 312.5\,\mathrm{kHz}$ and decimation factor $N = 312$. (a) In the time domain, the FIR filter has a length of $3N$ samples, and an apparent time delay of $3N/2$ samples. It is a convolution of three "boxcars". The area under the curve is 1. (b) In the frequency domain, notches occur in the amplitude spectrum at multiples of $f_{\mathrm{OUT}} = f_{\mathrm{MOD}}/N$. Only the first 6 lobes of the filter are shown; the Nyquist frequency is actually at 156.25 kHz. Once the time series is decimated to f_{OUT}, the Nyquist is reduced to $f_{\mathrm{OUT}}/2 = 500.8\,\mathrm{Hz}$. All energy in the signal above this point will be aliased.

An N-sample running average filter in Z-transform notation is

$$f_1(z) = \frac{1}{N} \sum_{k=0}^{N-1} z^k. \tag{4}$$

An equivalent expression for $f_1(z)$ is

$$f_1(z) = \frac{1}{N} \frac{1 - z^N}{1 - z}. \tag{5}$$

Equation 5 implies an efficient way to implement the filter/decimation step: (a) integrate with $1/(1-z)$; (b) decimate by N; and (c) differentiate the decimated series with $(1 - z')$ where $z' = z^N$ is the appropriate z-value to use on the decimated series. The advantage of the filter form in equation 5 is that multiple low-pass stages may be implemented with little more effort than that required for a single low-pass stage. The ADS1210 employs a three-stage running-average low-pass filter:

$$f_3(z) = \frac{1}{N^3} \frac{(1 - z^N)^3}{(1 - z)^3}. \tag{6}$$

Since filter f_3 is the convolution of three "boxcar" filters, it has a finite impulse response (figure 2(a)). The phase response of f_3 is linear, and is equivalent to a time delay of $3/2 \cdot \Delta t_{\mathrm{OUT}}$ where $\Delta t_{\mathrm{OUT}} \equiv 1/f_{\mathrm{OUT}}$ is the sample interval of the output time series. The amplitude spectrum of f_3,

$$|f_3(\omega)| = \left| \frac{\sin(N\omega/2f_{\mathrm{MOD}})}{\sin(\omega/2f_{\mathrm{MOD}})} \right|^3, \tag{7}$$

is shown in figure 2.

Experimental Response to a Square Wave Input

In order to verify that filter f_3 (equation 7) characterizes the amplitude response of the ADS1210 to an analog signal, a precision square wave was presented to the A/D converter and the digital output analyzed.

Complicating the analysis is the fact that the filtering and decimation steps are combined, with the ADS1210 issuing samples at the output frequency f_{OUT}. High frequencies in the input square wave (not insignificant due to the sharp edge present on the square wave) are folded about the Nyquist frequency $f_{\mathrm{OUT}}/2$ into low frequencies by the decimation process. For the purpose of measuring the unaliased filter response, it is desirable to recover the digital signal at the internal sampling rate f_{MOD} instead of the output sampling rate f_{OUT}. This can be done by employing a principle similar to that of a sampling oscilloscope: assume that the input is periodic, and measure the output over many cycles. Follow this with the construction of a single composite cycle by wrapping the output about the signal period. The composite cycle thus appears to be sampled at an arbitrarily fine rate provided that the signal period is *not* an even multiple of the output sample interval Δt_{OUT}.

Let us illustrate the formation of a finely sampled composite cycle with numbers from the actual experimental setup. The PAR4CH unit provides a 10 MHz clock signal to the ADS1210. With the Programmable Gain and Turbo Mode Rate settings of the ADS1210 at their default values of 1 and 16, respectively, the modulator sampling rate is

$$f_{\mathrm{MOD}} = 10\,\mathrm{MHz} \,/\, 32 = 312.5\,\mathrm{kHz},$$

and the modulator sampling interval is

$$\Delta t_{\mathrm{MOD}} = 1/f_{\mathrm{MOD}} = 3.2\,\mu s.$$

An output sampling rate of 1 kHz is desired, but since f_{OUT} depends on the modulator rate f_{MOD} and the oversampling factor N (equation 3), the best we can do is to set N to:

$$N = 312 \text{ (oversampling rate)},$$

which defines f_{OUT} and Δt_{OUT} to be:

$$f_{\mathrm{OUT}} = f_{\mathrm{MOD}}/N = 1.0016\,\mathrm{kHz} \text{ (approximate)},$$
$$\Delta t_{\mathrm{OUT}} = \Delta t_{\mathrm{MOD}} \cdot N = 0.9984\,\mathrm{ms} \text{ (exact)}.$$

The analog test signal applied to the PAR4CH was a square wave derived from a crystal-controlled oscillator. The oscillator frequency of 1.000 MHz was divided by 2^{14} to yield a square wave frequency f_{SQ} and period T_{SQ} of:

$$f_{\mathrm{SQ}} = 1\,\mathrm{MHz} \,/\, 16384 = 61.035\,\mathrm{Hz} \text{ (approximate)},$$
$$T_{\mathrm{SQ}} = 1\mu s \cdot 16384 = 16.384\,\mathrm{ms} \text{ (exact)}.$$

The two regulated voltage levels of the square wave were 0V and +5V. A Schmitt trigger was used to sharpen up the leading and trailing edges of the square pulses: the rise and fall time on each edge is approximately 20 ns.

The number of output samples in one cycle of the input square wave, $T_{\mathrm{SQ}}/\Delta t_{\mathrm{OUT}}$, is approximately 16.41, meaning that after every 16 samples, another cycle of the square wave is presented to the A/D converter apparently shifted to the right by 0.41 samples. Since 41 and 100 do not share common factors, the composite waveform assembled from 100 cycles of the input square wave appears to be sampled at a rate 100 times the actual sampling rate f_{OUT}.

In reality, the "exact" times T_{SQ} and Δt_{OUT} are precise only to the range of a microsecond. The 1 MHz oscillator generating the analog square wave has a timing accuracy of 0.1 percent. Temperature coefficient effects in both the test signal generator and in the PAR4CH were observed to cause timing drifts on the order of $1\,\mu s$ in 100 seconds. Consequently, there is enough randomness to the ratio between T_{SQ} and Δt_{OUT} such that an arbitrary number of cycles may be folded into the composite waveform, resulting in good uniform coverage at an arbitrarily fine sample interval.

Figure 3 shows an example of a composite waveform assembled from 200 seconds of digital output from a PAR4CH unit, representing about 12 000 cycles of the input square wave. Approximately 200 000 individual samples make up the waveform. Figures 3(b) and 3(c) show enlargements of two regions of the plot in figure 3(a). The individual samples are plotted as points, and although there is a random element to their spacing in time, the coverage is uniform enough to provide an effective sampling interval of 100 ns.

The high resolution in time of the composite waveform is confirmed by the "stair-step" features on the waveform seen in figures 3(b) and 3(c). The unexpected features are a result of the sample-and-hold circuitry on the ADS1210. Each modulator sample interval, spanning $3.2\,\mu s$, consists of a sample window followed by a hold window. During the sample window, a low-impedance path from the input signal to the sampling capacitor allows the capacitor to charge to the input voltage level (+5 V). During the hold window, the charge on the capacitor is presented to the comparator. On subsequent cycles, the leading edge of the square wave is progressively delayed, and the sampling capacitor has progressively less time to charge from 0 V to +5 V during the time the sample window is open. Consequently, the sample points in one "step" in figure 3(b) effectively trace out the exponential RC charge curve on the sampling capacitor. The reported RC time constant of the ADS1210 sample-and-hold circuitry [2] is $2\,\mu s$ (i.e. $8\,\mathrm{pF} \cdot 250\,\mathrm{k}\Omega$) which is consistent with the charging curve seen on each step in the figure.

Ideal versus Measured Filter Response

For the subsequent filter analysis, the composite waveform of figure 3 was resampled to the modulator rate f_{MOD}, with samples taken from the "hold" window in the modulator sample interval, that is, on the flat portion of each stairstep in figure 3(b).

The amplitude spectrum of the waveform from figure 3 is drawn in black in figure 4(a). This represents the spectrum of a square wave filtered by the ADS1210 for frequencies all the way up to the modulator Nyquist $f_{\mathrm{MOD}}/2$. Figure 4(b) shows the spectrum for frequencies that are more in the range of interest, up to 10 kHz. In fact, the meaningful (unaliased) frequencies of the output are no greater than the output Nyquist, $f_{\mathrm{OUT}}/2 = 0.5\,\mathrm{kHz}$. The spectral holes are positioned at multiples of f_{OUT}, as expected.

The gray curve in figure 4(b) is the ideal spectrum of a pure square wave with the same power and period as the experimentally-generated square wave. The amplitude response of the low-pass filter is derived by dividing the black curve by the gray curve, or (in the log-amplitude domain) by taking the difference of the two curves. The resultant measured spectrum of the filter is shown in figure 5. For comparison, the theoretical amplitude response of the low-pass filter (equation 7) is drawn in gray on figure 5. There is a good match between the measured and the theoretical amplitude spectra out to the 6th aliased lobe (6.5 kHz).

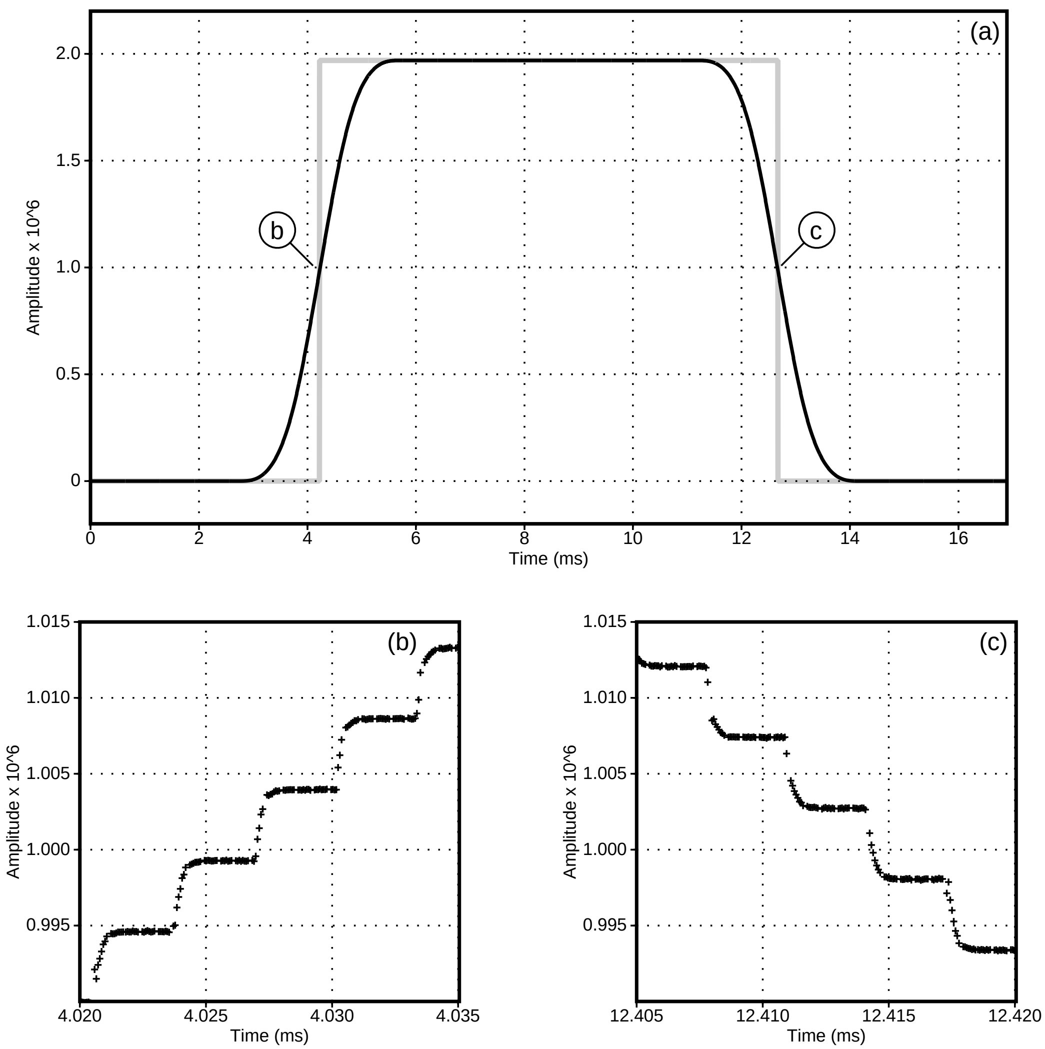

Figure 3: Response in the time domain to a periodic square wave. The black curve in (a) shows one cycle of the filtered square wave, with a period of 16.384 ms. The curve is a "stack" (composite) of 12 000 cycles from the output of the PAR4CH. As a consequence, the sampling is irregular, with an average sample interval of approximately 100 ns. The PAR4CH unit puts out samples as 24-bit integers; an amplitude of 0 on the plot represents 0 V on the differential input, and an amplitude of approximately 1 950 000 represents an input of 5 V. For comparison, the ideal input square wave is shown in gray, scaled to match the measured curve. The actual edge rise and fall times of the input square wave are on the order of 20 ns. Two regions of the black curve are shown magnified in (b) and (c), in which individual samples are plotted. The stairstep appearance of the curve is a consequence of the sample-and-hold circuitry in the modulator: every 3.2 μs, the signal voltage is presented to the sampling capacitor and held for the comparator to make its choice. Each stairstep resembles the shape of an RC charge (and discharge) curve.

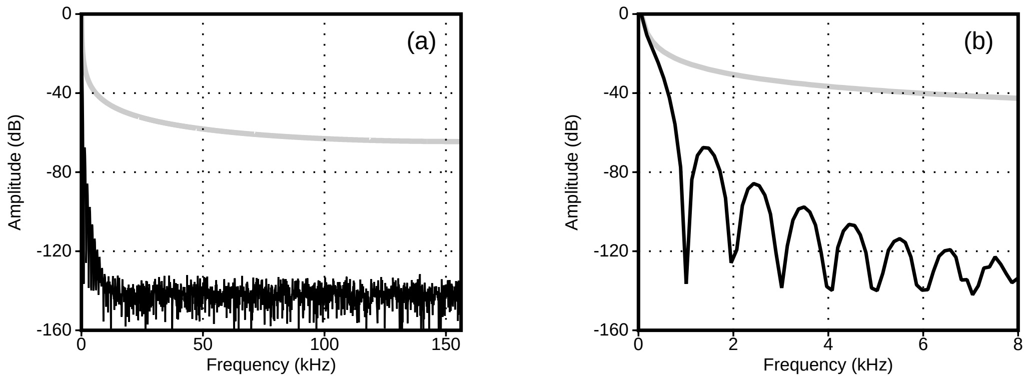

Figure 4: Amplitude spectrum of filtered square wave. Plot (a) shows the full frequency range to the Nyquist $f_{\mathrm{MOD}}/2$. Plot (b) shows the same curve limited to 10 kHz. Beyond 10 kHz the spectral components drop below the -140 dB floor representing sample roundoff noise (the samples in the time series are stored in 32-bit IEEE floats). The gray curve is the amplitude spectrum of the ideal, unfiltered square wave. Since the even-indexed frequency components of a periodic square wave are zero, only the odd-indexed frequency components are shown.

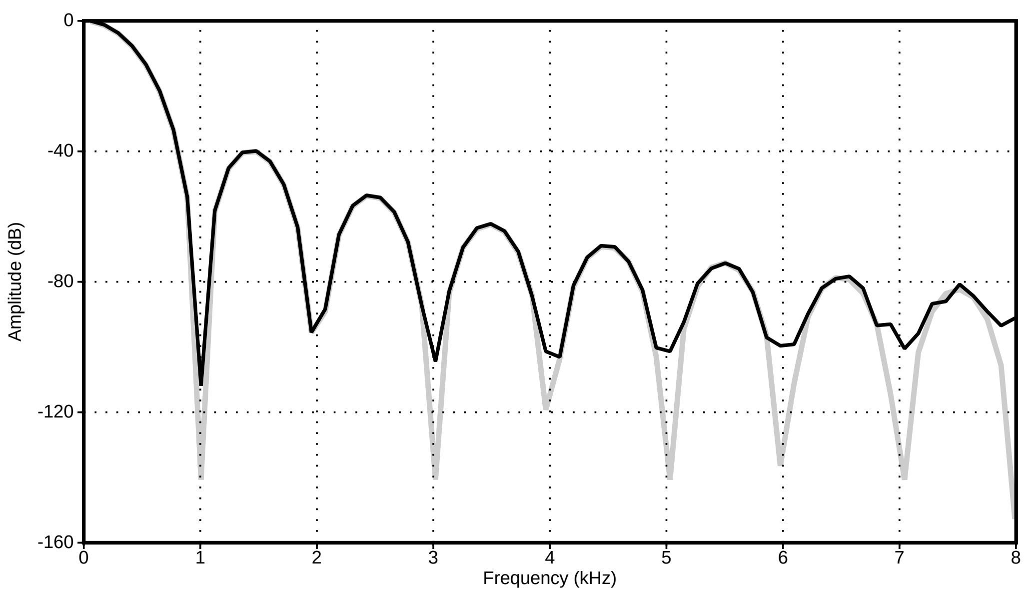

Figure 5: Comparison between the theoretical and measured filter spectra. The black curve is the measured amplitude spectrum of the 3-stage low-pass filter, estimated from a square wave input. It is the difference between the two curves in figure 4. The gray curve is the theoretical amplitude spectrum from equation 7, sampled at the same rate as the black curve. The spectral holes at 1 kHz intervals are more pronounced if the sampling of the frequency axis were made finer.

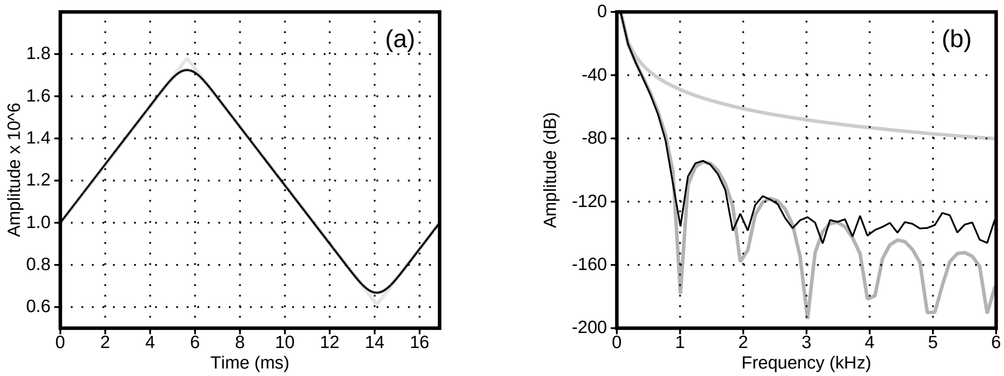

Figure 6: Response to a triangular wave input. Three curves are shown: black indicates the measured response, gray indicates the theoretical filtered response, and light gray indicates the theoretical unfiltered input signal. (a) In the time domain, the measured and theoretical curves overlap precisely. (b) In the frequency domain, the measured and theoretical curves diverge at the point where the limit in sample precision is reached: about -140 dB or 23 bits.

Triangular Wave Input

Similar tests were performed with a periodic triangular wave signal as input. Figure 6 shows the results of such tests. The black curve in figure 6(a) illustrates how the low-pass filter smooths the corners of the triangular wave as expected. Both theoretical and measured curves are plotted in figure 6(a), but the match between the two is so close at this scale that the overlap is exact. A third curve, the unfiltered triangular wave signal, is also shown in light gray.

The amplitude spectra of the three curves in figure 6(a) are plotted in figure 6(b). As in the case of the square wave input signal, the frequency components of the measured triangular wave match those of the theoretical filtered wave as long as the component amplitudes are above -140 dB, the roundoff noise floor. In the triangular case, the noise floor at -140 dB is reached sooner (by approximately 3 kHz), since the theoretical rate of decay of the frequency components is twice that for the case of the square wave.

Antialias Filter Issues

The fact that, in sigma-delta A/D converters, the input signal is oversampled by the modulator and digitally filtered does not alone prevent aliased energy from corrupting the output time series. Take for example a desired sample rate of 100 Hz. From figure 5 it is seen that signals from the Nyquist (50 Hz) up to about 80 Hz are passed with a range of attenuation from 10 dB to 40 dB. At the 80 dB attenuation level, signals up to 200 Hz may be passed. Placing an analog antialias filter in front of the PARxCH unit is one obvious solution, but problems common to analog filters may prove unacceptable, such as the resulting distortion of the phase spectrum.

One alternative to an analog prefilter is to move part of the task of oversampling/filtering from the A/D converter downstream into the computer. For example, if a final sample rate of 100 Hz is desired, the PAR4CH may be configured to output samples at 1 kHz. The subsequent filtering and subsampling to 100 Hz may be done on the computer, with the freedom of being able to apply a custom-designed digital antialias filter with zero phase distortion. The antialias filtering stage may also restore the precision to the final 100 Hz samples that may have been lost from sampling at the higher rate of 1 kHz in the PAR4CH unit.

Conclusion

We have experimentally verified that the amplitude response of the built-in low-pass filter in the PARxCH A/D converters as given in equation 7 is correct. Figure 5 illustrates the comparison between the theoretical amplitude response and that calculated from a square wave input signal.

References

[1] Park, S., *Principles of sigma-delta modulation for analog-to-digital converters*, Motorola application note APR8/D, revision 1.

[2] Burr-Brown, *ADS1210/1211 24-bit analog-to-digital converter data sheet*, Texas Instruments (2000).

[3] Baker, B., *How to get 23 bits of effective resolution from your 24-bit converter*, Burr-Brown application note AB-120.