OVER QUANTIZING, THE NEXT STEP IN DIGITAL SIGNAL PROCESSING

William T. Hicks, P.E.
L-3 Communications, Telemetry East

ABSTRACT

Simplification of the analog front end of a signal conditioning circuit can be accomplished by over-quantizing the input signal and using DSP for gain and offset. In this case, a much higher precision A/D converter is used than required by the desired output accuracy. The excess bits are then used to allow the DSP math to give an effective gain to the signal. By a similar function, offset of over 100% can be mathematically removed as long as the input signal does not exceed the A/D input voltage range.

KEY WORDS

Over-quantizing, DSP, A/D, gain, offset.

INTRODUCTION

Individual analog channels of airborne telemetry systems require specific gains and offsets that are unique to the transducer that is being monitored. Traditionally these channels have gains and offsets that are either fixed or selected from a limited table of values. In addition there are inaccuracies in the components used. The result is that the gains and offsets often have errors that exceed the output resolution of one least significant bit (lsb). This then requires that the channels be calibrated. In the past this has been accomplished by mounting fixed resistors in a select-by-test method.

The method of gain, offset, and calibration presented in this paper has been made practical by the availability of faster, smaller, and more accurate analog to digital (A/D) converters and the advent of inexpensive and small digital signal processors (DSP).

BACKGROUND

This program was initiated to meet some new requirements in a small custom airborne telemetry system with about 200 analog channels, each requiring gain, offset, and anti-aliasing filtering. In order to meet the space requirements, it was decided to over-sample the channels and to take advantage of DSP to do the filtering. This would simplify the input circuits by relaxing the filtering required in the analog front end, in both the number and the accuracy of components.

The next problem was the potential variability of the gain and offset of each channel, with the final values not known until after the system would be built. While a gain range was known, the specific gain value was not known. If the worst case anticipated input signal level was used to set the gain, then the resolution of each bit would be reduced, making a 12 bit system look more like 8 bits. This could have been solved with a select-by-test component at each front end, but the final values might not be known until after assembly of the boxes, causing expensive field changes. Another option considered was a variable gain front end or mid stage amplifier. Because of the size constraints, this option was not viable.

In the past we had made units with fine gain and offset adjust (a few percent) in the DSP portion, and this could be used for calibrating out component tolerances. This method causes no loss in overall system accuracy, as long as the change is small. We had also used A/D's with an extra bit to use for octave offsets. This was done by setting up a +/- 5 volt, 11 bit A/D to be used with either +/- 2.5 volt inputs or 0 to +5 volt inputs, and then using logic to select the needed output bits in a 10 bit format. Because we were using 10 bits of an 11 bit converter, we maintained the required 10 bit accuracy.

NEW METHOD

The new method of conversion and processing takes this technique to the next logical step. While the system requirement is for 12 bits, it is now easy to find 16 bit A/D's and references to drive them, that maintain 16 bit accuracy over an extended temperature range. These new parts are small in size and low in power, while allowing over-sampling of the input signal for DSP filtering. This usage of a 16 bit converter when only 12 bits are needed gives 4 bits of excess resolution. This excess resolution can be used to get four octaves of gain or offset. Now the input amplifier only has to be set up with a fixed gain that will bring the worst case maximum signal level in below the A/D clipping level. The 16 bit converter will resolve the input signal into 64,000 parts when only 4000 parts are needed. The DSP can now be programmed to use the needed portion of the converter that contains the signal. Because this can be programmed into the unit, it can be changed easily at any time, as long as the design allows reprogramming of the unit after assembly.

In addition to this, the memory of the DSP can be used with multiple look-up tables for calibration. One table can be loaded when the unit is calibrated at the factory. This table removes any errors in the analog components or A/D converter. The box can now be calibrated so that it can be interchanged with other boxes and no changes in output level are observed. A second table contains the user transducer calibration information. This information is loaded with the sampling format, and as a box is replaced, this transducer calibration information loads up with the format put in to the box by the user.

DETAILED DESCRIPTION

Figure 1 is a block diagram of one channel in the system. The input signal is first amplified (if required) by an input amplifier with a fixed gain. This feeds an anti-aliasing filter with sufficient

filtering at a high frequency to prevent aliasing at the over-sampling rate. This feeds a 16 bit A/D converter that has both DC and AC accuracy to 16 bits. The output of the A/D goes to the (DSP), which is a high speed math processor that uses Finite Impulse Response (FIR) filter techniques to filter the signal to any desired shape and response, by altering the coefficient table used in the calculations. At any point in time the coefficient table can be modified to change the filter characteristics. This can be accomplished by downloading new values over a standard computer port on the system, or during flight by selecting alternate formats that have been pre loaded. By using DSP and FIR filters, a much sharper filter is easily obtained. This can prevent ailiasing with a cutoff frequency versus sample rate of one to four. Also by using a DSP FIR filter, the accuracy of the filter varies only with the changes of the system crystal oscillator, which is much more accurate and stable then discrete analog filter components.

Figure 1
Block Diagram

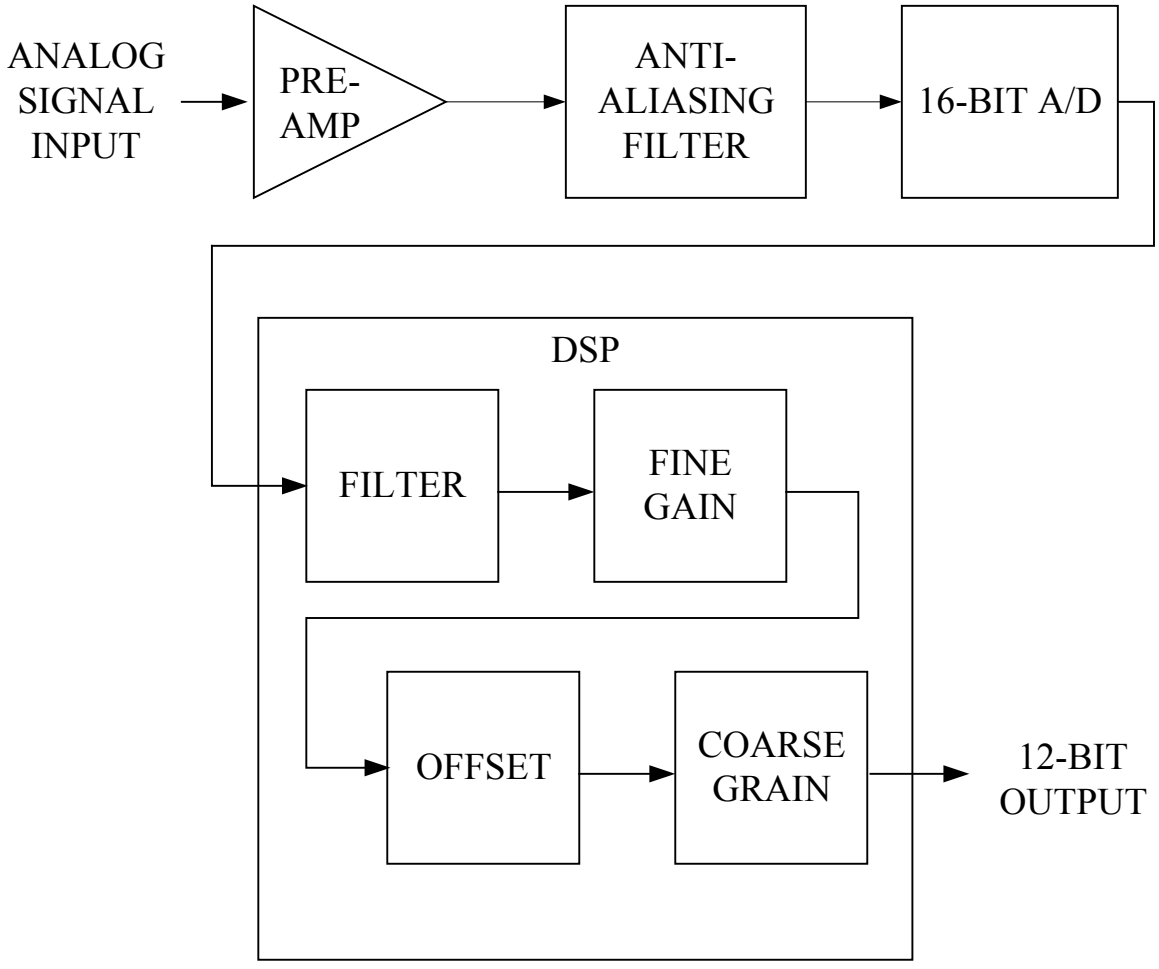

The sampling rate of the input signal is synchronized to the sample rate of the signal in the PCM stream. This prevents beat frequency artifacts from being produced by sampling a sampled signal at a non-synchronous rate.

After the filter stage is the is the fine gain stage. The fine gain is a fraction between one half and one. Because the gain is fractional, there is no possibility of saturation of the output. Next is the stage that inserts offset. This moves the signal as far as possible from the full scale. Again there is no chance of saturation. The last stage is the final gain stage, which is a course gain that is accomplished by shifting the input, which is used for octaves of gain. This brings the full scale input to be just within the DSP maximum available range.

With the ability to fine-change gain and offset, a utility program can be used by the end user to monitor the PCM output, and with a transducer in the calibration position, the program automatically corrects the calibration error and loads new gain/offset values in the format tables.

The usage of the DSP chip that is the focus of this paper is to do gain and offset adjustment. Because a 16-bit A/D is being used for a 12-bit system, 4 bits, or 24 dB, of excess gain is available. By properly scaling the input pre-amplifiers associated with each channel, this extra 24 dB can be used to multiply the output of the FIR calculations by a user-inserted value to fine-tune the gain. In addition, by proper positioning the input signal in the A/D input range, full scale offset can be accomplished. Again, there is up to 24 dB of offset available, however at the expense of gain adjustment.

Another consideration is that the DSP has the capability of doing engineering unit conversion and linearization. If this is desired, then ground processing is reduced. The DSP can also be used to do things like RMS averaging, frequency measurement, and FFT. These properties can allow the reduction of the signal list and its associated bandwidth, or conversely, increase the number of signals available in the existing bandwidth.

EXAMPLES

The technique is best described by giving some examples of input signal voltages and how the output values are calculated. Consider a system where the A/D input range is –10.24 to +10.24 volts and that the input amplifier has unity gain. The voltage range allows signals from –10 volts to +10 volts to be measured without worrying about small offsets. This gives a bit weight of 20,480,000uV/2^16 = 312.5 uV/bit. In the DSP all arithmetic is done in "1.15" signed fractions (sign bit, decimal point, 15 digit fraction). The output of the A/D is:

Input voltage (volts)	A/D code out (hex)
+10.24	7FFF
0	0000
-10.24	8000

The DSP FIR filter will multiply and accumulate the input by a set of coefficients that give unity gain when added together, therefore, outside of the transition band, they have no effect. (i.e.: a five tap filter might have coefficients of: 0.3, -0.1, 0.6, -0.1, 0.3; which sum to one). Note that the DSP multiplier-accumulator has a 40 bit accumulator register for two 16 bit words and rounding instructions for the final answer. This removes error accumulation during processing.

The output of the system is changed to 12 bit words which have the following weights.

Input	DSP code out (hex)
+full scale	7FF
0 scale	000
-full scale	800

Example 1.

The first case is a -8 volt to + 8 volt signal from an accelerometer. The user desired 12 bit output is 4 mV/bit, with a full scale output of +/- 7D0h counts (+/-2000 counts). Full scale in and out is shown. This case represents a signal that matches closely to the maximum levels of the A/D. The fine gain is 3E80/6400h. The output gets no offset. The coarse gain is two. The DSP output is the top 12 bits of the calculation. Figure 2 shows the input voltage and the progression of steps.

Figure 2
+/- 8 Volt Input Signal

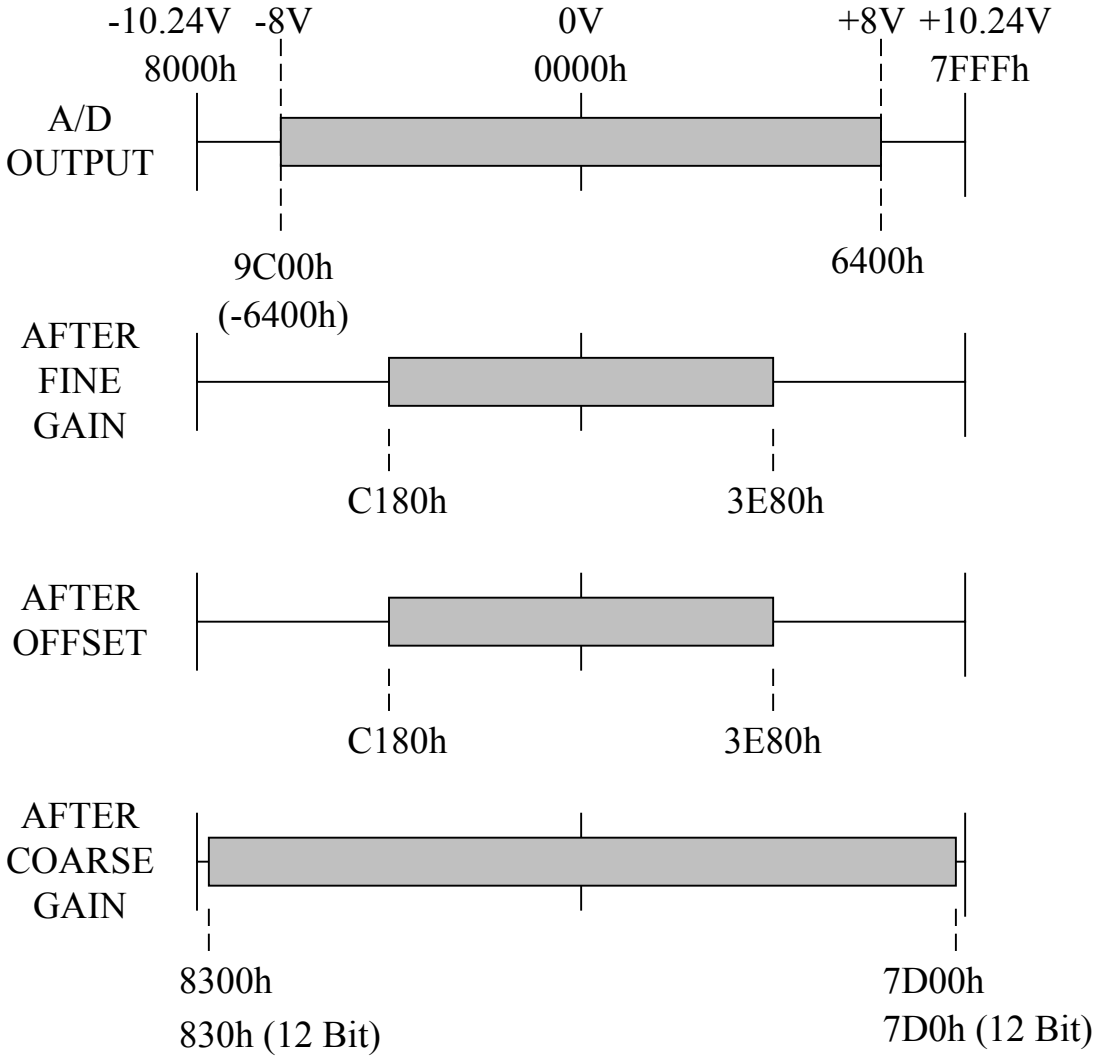

Example 2.

The next case is a –0.5 volt to + 0.5 volt signal from a strain gauge. The desired 12 bit output is 0.25 mV/bit, with a full scale output of +/- 7D0h counts. Full scale input and output is shown. This case represents a signal that matches closely to the minimum levels of the A/D. The fine gain is 3E80/6400h. The output gets no offset. The coarse gain is 32. The DSP output is the top 12 bits of the calculation. Figure 3 shows the input voltage and the progression of steps.

Figure 3
+/- 0.5 Volt Input Signal

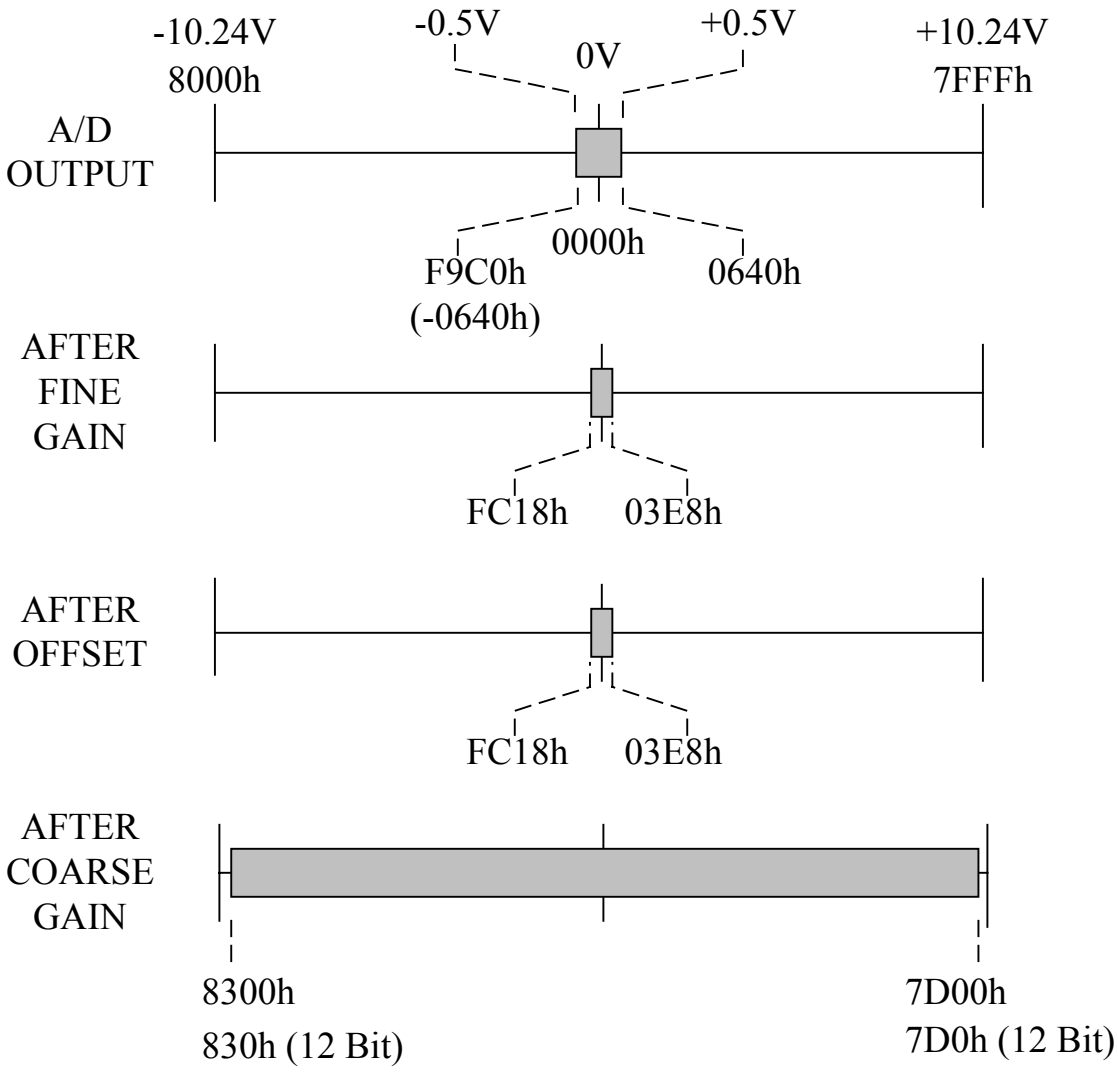

Example 3.

The next example is an RTD with a DC bias on the positive side. It has a low scale input of –50 degrees or 79.96 ohms and a 20 mA excitation or 1599.2 mV It has a full scale input of +350 degrees or 231.89 ohms and a 20 mA excitation or 4637.8 mV. Both DSP offset and gain will be used. The fine gain is 1F40/25FCh. The output gets 2012h offset. The coarse gain is eight. The DSP output is the top 12 bits of the calculation. Figure 4 shows the input voltage and the progression of steps.

Figure 4
+1.5992 TO +4.6378 VOLT INPUT SIGNAL

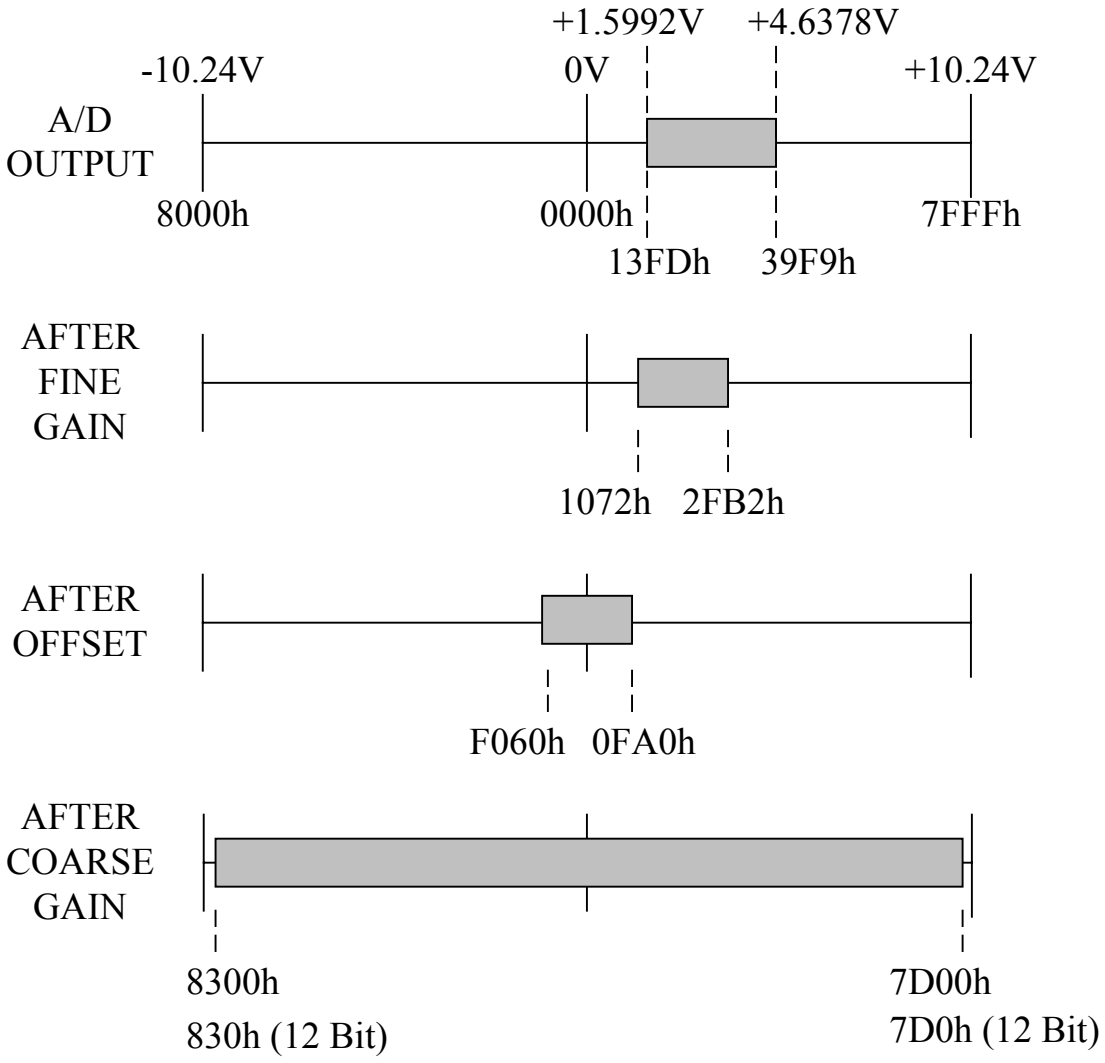

Example 4.

The next example is the same RTD with a DC bias on the positive side. It has been found that the high temperature is actually +450 degrees. Therefore it has a low scale input of –50 degrees or 79.96 ohms and a 20 mA excitation or 1599.2 mV It now has a full scale input of +450 degrees or 266.94 ohms and a 20 mA excitation or 5338.8 mV Both DSP offset and gain will be used. The fine gain is 1F40/2EBFh. The output gets 1CFDh offset. The coarse gain is eight. The DSP output is the top 12 bits of the calculation. Figure 5 shows the input voltage and the progression of steps.

Figure 5
+1.5992 TO +5.3388 VOLT INPUT SIGNAL

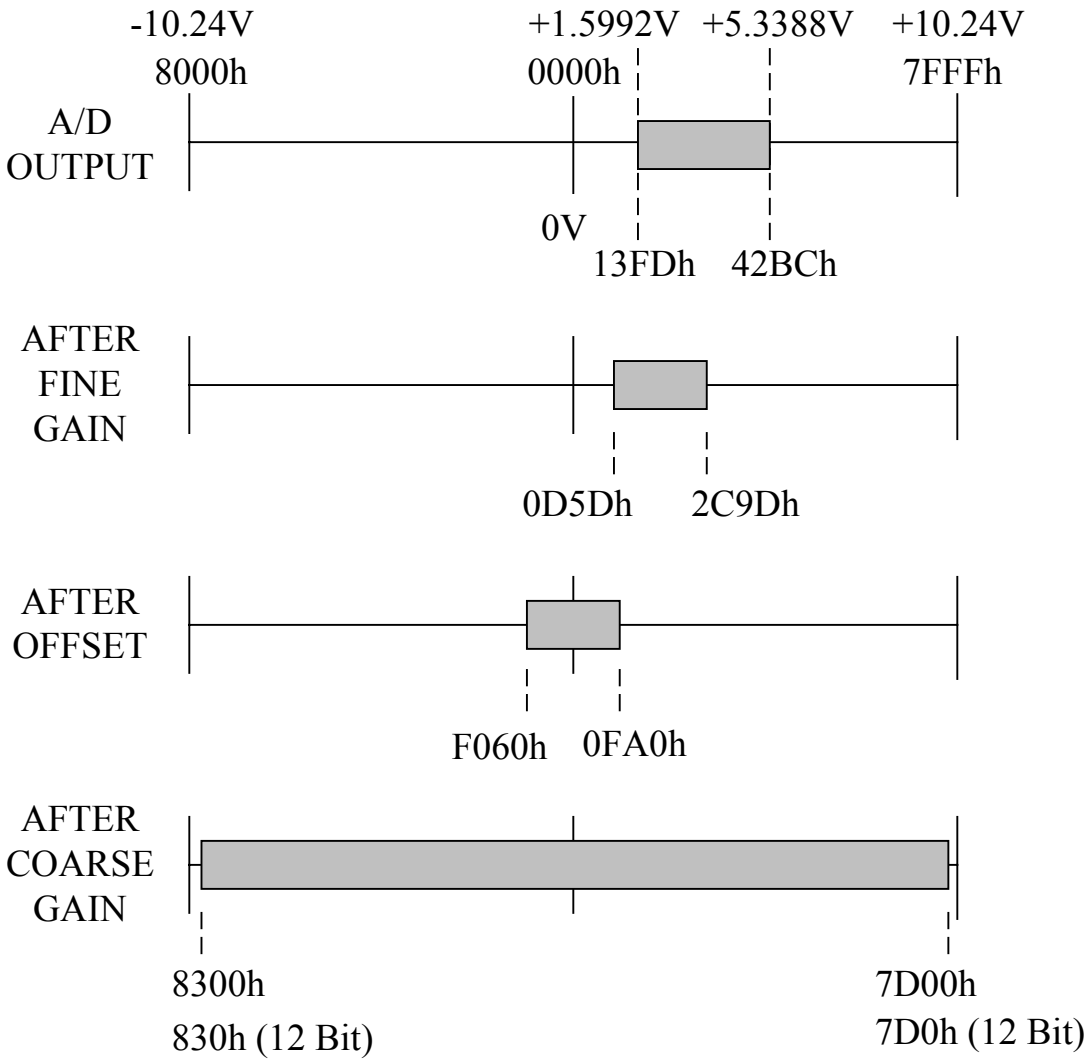

CONCLUSION

The examples give the best explanation of how the new method can be used to simplify the hardware needed to have a variable gain/offset system. These examples show how changes are simplified by minor adjustments to the software parameters of the DSP processor. This allows a maximum of flexibility at a minimum of cost.