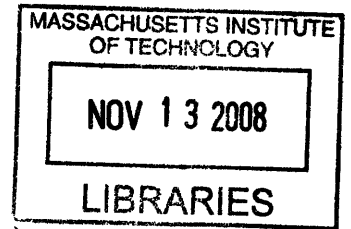


**Increasing A/D Conversion Resolution by
Dynamic Scale Adjustment**

by

James J. Wnorowski

S.B., Massachusetts Institute of Technology (2006)



Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of
Masters of Engineering in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2008

© Massachusetts Institute of Technology 2008. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
September 8, 2008

Certified by
Gerald Jay Sussman
Panasonic Professor of Electrical Engineering
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Department Committee on Graduate Theses

ARCHIVES

Increasing A/D Conversion Resolution by Dynamic Scale Adjustment

by

James J. Wnorowski

Submitted to the Department of Electrical Engineering and Computer Science
on September 8, 2008, in partial fulfillment of the
requirements for the degree of
Masters of Engineering in Computer Science and Engineering

Abstract

Dual-Ranging analog-to-digital systems create performance gains by exploiting the dynamic range of certain classes of signals. This thesis explores the improvements possible when using a dual-ranging system to gain resolution when digitizing audio waveforms. Experiments indicate between 3 to 6 *dB* of possible improvement, results comparable to gains attainable through other techniques.

Thesis Supervisor: Gerald Jay Sussman

Title: Panasonic Professor of Electrical Engineering

Acknowledgments

I would like to thank the following groups and people:

Vtech Engineering Corporation, for their willingness to take me on as a part-time employee and let me take my thesis data on the company's data acquisition board... provided I helped develop firmware for it, of course. Special thanks go to Paul Butler, who guided me through VHDL and DSP C, as well as helped me debug hardware issues well into the evening as I attempted to take my thesis data.

The Undergraduate Course VI Office, for always being willing to work with me to get through the the requirements of the curriculum, even when I found out I had attempted to double count a course during my last undergraduate term. Special thanks to Anne Hunter for inviting me to come in and talk about anything (or nothing) whenever I needed to unwind for a bit.

Mandi Holmes, for spending the better part of two days putting this document into L^AT_EX. She made sure I was thinking about my thesis, even when it was the last thing on my mind.

Professor Sussman, for agreeing to supervise my M.Eng work after I'd spent a year quietly ignoring the thesis requirement. His broad base of knowledge and interest in discussing nearly any subject inspired me once more to learn for the sake of learning.

My parents, for without all their help and support I would certainly not be where I am today. As far as I can tell, good parenting is a balancing act between guiding a child through life while letting him chart his own course. I believe they got it more or less right.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1	Introduction	11
2	Background	13
2.1	Basics of Analog to Digital Converters	13
2.2	The Successive Approximation A/D Converter	14
2.3	Errors in A/D Converters	15
2.4	Fundamental Performance Limits	16
2.5	Other Resolution Enhancing Techniques	16
3	Dual Range A/D Conversion	19
4	Implementation	25
5	Results	29
6	Conclusion	33
A	Sample MATLAB Script	35

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

2-1	Block Diagram of Successive Approximation A/D Converter [2]	14
3-1	Bach, Partita in C minor	20
3-2	Thuille, Piano Sextet in B flat major	20
3-3	Theme Song from <i>Friends</i>	21
3-4	The Dual-Range A/D Conversion System	23
4-1	Dual Channel Conversion Experimental Setup	27
5-1	Curve Fitting: No Distortion	31
5-2	Curve Fitting: Moderate Distortion	31
5-3	Curve Fitting: Significant Distortion	32

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

Since the first commercial analog-to-digital converter was released in 1954, analog-to-digital conversion has been progressing for the past half a century [7]. As is the general trend with technology, new processes have provided incremental improvements while novel designs have created different entirely new classes of converters. These different families of converters fill different design requirements: some run at high speed but with low accuracy, others run slower but with much higher accuracy. However, process gains are approaching physical limits inherent in integrated circuits. What's more, the resolution of state of the art A/D converters is reaching a fundamental physical limit: that of thermal noise.

One can attempt to obtain additional resolution by applying an algorithm to the data to be collected. One such algorithm would be to dynamically allocate the range of the converter, allowing the converter to become more precise for low amplitude signals. Of course, this method requires that the data have wide dynamic range.

While ideal, this method has serious challenges in implementation. Another method to implement this algorithm is to operate two or more converters in parallel with different scales, and use digital logic to pick the converter that is best suited to the current value of the signal. This is the approach that will be followed in this research.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Background

This chapter will cover some background information on analog-to-digital conversion. Topics include the basics of analog-to-digital conversion, analog-to-digital performance measures such as signal to noise ratio, Johnson noise as a limiting case of analog-to-digital resolution, and other techniques commonly used to enhance analog-to-digital performance.

2.1 Basics of Analog to Digital Converters

An analog-to-digital converter (A/D) is a device that takes an analog input, samples it at a certain rate, and quantizes the sampled input into a digital output. A/D converters are typically represented by two important parameters: the resolution and the sampling rate. The resolution determines the number of bits in each sample output. The sampling rate determines the maximum bandwidth of the sampled signal, due to the Nyquist-Shannon sampling theorem.

Analog-to-digital converters can be unipolar, where the allowable input voltages range between zero and some maximum voltage (called the full scale voltage and denoted V_{fs}), or they can be bipolar, in which case the allowable input range is between two voltages of equal magnitude but opposite sign (usually denoted $+\frac{V_{fs}}{2}$ and $-\frac{V_{fs}}{2}$).

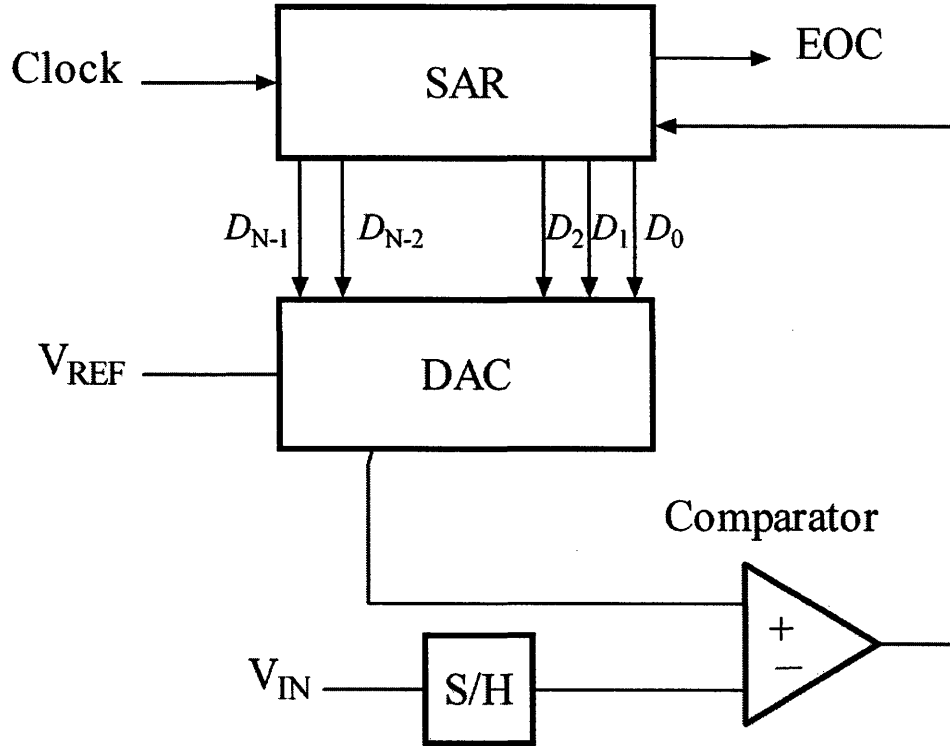


Figure 2-1: Block Diagram of Successive Approximation A/D Converter [2]

2.2 The Successive Approximation A/D Converter

There are several different types of A/D converters available. Each architecture type determines the properties of the converter, such as conversion speed and conversion resolution. One such architecture, the Successive Approximation ADC, is explained in Figure 2-1.

Figure 2-1 breaks down the successive approximation converter into four sub-circuits: a Successive Approximation Register (SAR), a Digital to Analog Converter (DAC), a Comparator, and a Sample and Hold (S/H). At the start of a conversion initiated by the clock, the S/H samples the voltage input, and holds that voltage for the duration of the conversion. The comparator then compares the input voltage to the voltage generated by the DAC.

The first voltage generated by the DAC is half converter scale: the SAR sets D_{n-1} high (to 1), while all the other bits are set to zero. If the comparator determines the input voltage was greater than the DAC voltage, the SAR holds the D_{n-1} bit. If,

however, the input voltage was less than the DAC voltage, the SAR resets the D_{n-1} bit to 0. In either case, the D_{n-2} bit is set to zero and the process continues. Thus, the successive approximation algorithm is a binary search tree designed to gain an extra bit of resolution at each converter step. Once all N bits have been set, the SAR asserts the End of Conversion signal, indicating that the data is ready to be read, and another sample can be taken.

2.3 Errors in A/D Converters

There are two general types of errors present in analog-to-digital converters: dc errors and ac errors. There are 4 types of dc errors: offset error, gain error, and two types of linearity error. In general, the offset error and gain error can be trimmed by the user: the offset is trimmed at zero input, and the gain error is trimmed to the converter's full scale. For ac applications, such as digital audio, the dc specifications of the data converter are not as important as the ac specifications, however.[6]

The ac specifications of an ideal A/D converter arise from the sampling and quantization processes. The maximum error an ideal converter makes for a particular dc signal (voltage) is $\pm\frac{1}{2}$ LSB. If the signal is of large enough amplitude and aperiodic with respect to the sampling frequency, the quantization error will be uniformly distributed between these two values. This results in quantization noise with an rms value equal to the standard deviation of the normal distribution: $\frac{q}{\sqrt{12}}$, where q is the weight of the LSB. For this ideal converter, the Signal to Noise ratio is $6.02N+1.76$ dB, where N is the number of bits in the converter.

Of course, the noise and distortion of a practical converter will be worse than this theoretical ratio. However, this calculation can be done in reverse: the ratio of the RMS of the amplitude of the signal applied to mean noise and distortion present (SINAD, or Signal-to-Noise-and-Distortion Ratio) can be found through measurement. Using SINAD in place of SNR, the equation can be rewritten to solve for N .

$$N = \frac{SINAD - 1.76}{6.02}$$

This value is called the Effective Number of Bits, or ENOB. It is a simple measure of the ac accuracy of the converter. Today's cutting-edge 24 bit sigma-delta ADCs typically have an ENOB of around 20, indicating that the high twenty bits contain actual signal content, while the low four bits have likely been lost to noise.

2.4 Fundamental Performance Limits

Analog-to-Digital converters have been constantly improving over the years, as new A/D topologies have been invented and integrated circuit fabrication processes have become more advanced. Of course, there are hard limits to technologies where further improvement is physically impossible. For measuring the accuracy of A/D converters, that limit is called Johnson noise, or Thermal Noise. At room temperature (300K), for a bandwidth of 10kHz and assuming a 50 ohm load resistance, the power of the noise is on the order of $-134dBm$. This is on par with the quantization noise of an ideal 22 bit A/D.

2.5 Other Resolution Enhancing Techniques

Even though this theoretical performance cap has not yet been reached, historically it has been the case that certain applications have demanded performance in excess of what the data converters of the day were able to achieve. This need drove users to develop novel techniques to increase performance of their data acquisition systems. Often, these techniques involve using multiple converters to acquire data in parallel, and process the data in a clever fashion. Two such techniques are commonly used. Signal averaging is employed to gain resolution without trading speed, while time interleaving is used to increase the sample rate without sacrificing resolution.

The idea behind signal averaging is to have a number of ADCs converting the same signal, and then digitally sum the outputs. Each ADC's output will have a signal term and a noise term. The signal terms will all sum directly. The noise terms, however, should be uncorrelated, and therefore sum as the square root of the sum of squares. Carrying out the math, this process results in a $3dB$ increase in the SNR per doubling of converters. Thus, a system summing the outputs of four identical converters has an ENOB of one more than a system with only one such converter.

Time interleaving works to increase the sample rate of a converter system M times by using M converters. These M converters are evenly spaced in phase; e.g. 4 converters would be phased 90 degrees apart. In this manner, the analog sampling bandwidth is multiplied. Time interleaving can be used to increase the SNR as well. If a system's noise is uniformly distributed in frequency, time interleaving will spread that noise throughout the new bandwidth. A low-pass filter can then eliminate the noise outside the frequency band of interest. With a four converter system, time interleaving can produce the same $6dB$ of SNR improvement that signal averaging provided. However, great care must be taken to match the converters' gains, phases, and offsets, as each of these provides a source of distortion. This makes time interleaving more difficult to implement than signal averaging.

Both of the above methods are able to increase system performance in most cases. It should be noted that they each make different assumptions about the characteristics of the signal. In the case of signal averaging, the stated performance gain assumes that the noise in each channel is uncorrelated. For the time interleaving method, it is assumed that the system noise is wideband and spread evenly throughout the spectrum. If either of these assumptions is incorrect, the corresponding method will fail to yield the expected performance gain.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Dual Range A/D Conversion

The two prior methods were able to achieve modest gains in converter performance (SINAD) with fairly general restrictions on the signal being digitized. If we are able to impose more specific requirements on the input signal, it should be possible to achieve greater performance gains.

One possible requirement to impose on a signal is an amplitude constraint. A signal whose amplitude only reaches half the full scale value for a unipolar converter (or between $-\frac{1}{4}FS$ and $+\frac{1}{4}FS$ for a bipolar converter) will not exercise the MSB of that converter, and the actual resolution attained will be one bit less than the resolution available from the converter. Therefore, it is obvious that the peak value of a signal should match the full scale range of a converter to achieve best performance.

It may be the case, however, that these peaks are few and far between, while the average value of the signal lies well below converter full scale. This could be the case when the signal is a particular type of audio signal. For example, investigating *The OnClassical Collection* (a collection of classical music chosen by the author due to its distribution under a Creative Commons license), the majority of classical music had very infrequent peaks and average values under ± 0.5 , the half scale value for this medium.¹ (The data set was composed of WAV files processed and graphed in

¹It should be noted that not all music fits this pattern. An article in the IEEE spectrum [10] show that current popular music has greatly reduced dynamic range when compared to music from a couple of decades ago. This trend, dubbed the Loudness Race, resulted in modern popular music becoming louder and louder in an effort to gather more interest from people listening to music over the radio or on television. See Figure 3-3 for an example of this effect.

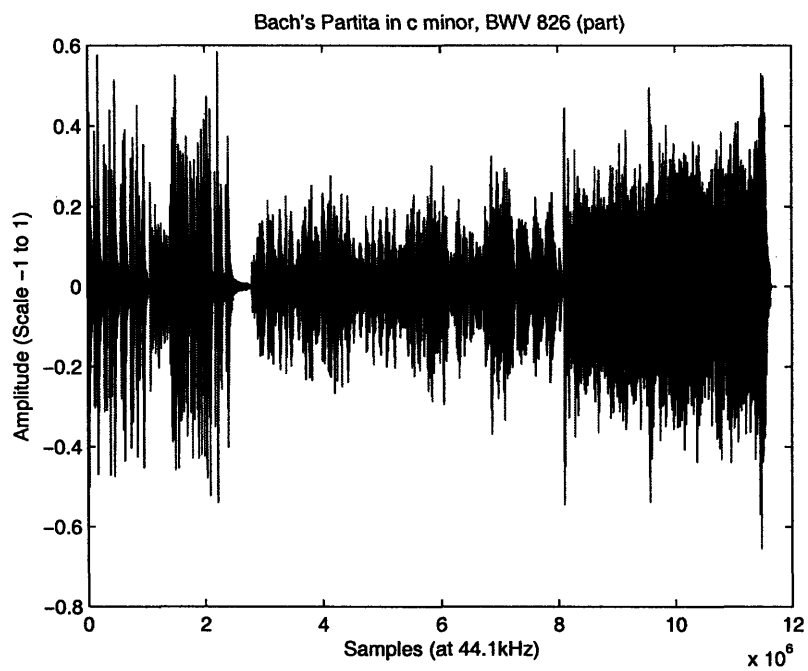


Figure 3-1: Bach, Partita in C minor

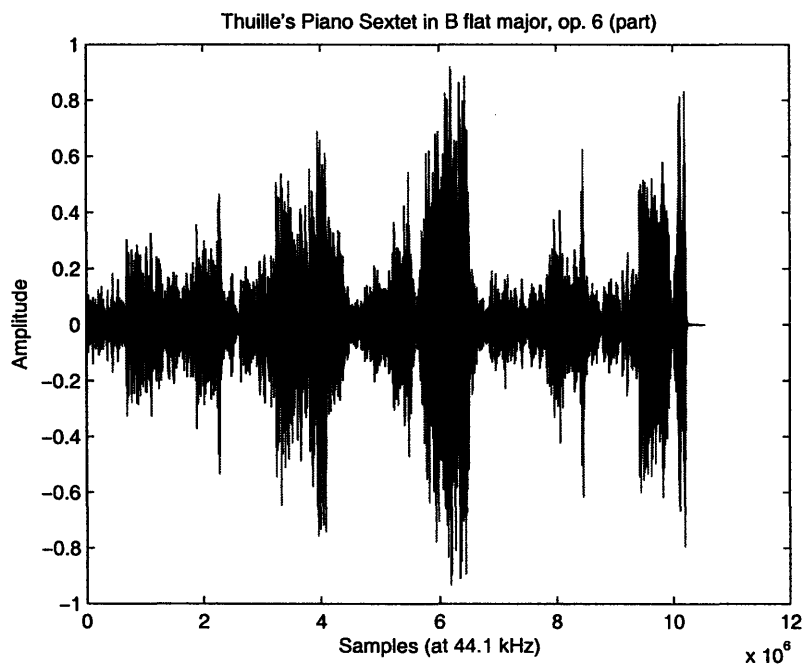


Figure 3-2: Thuille, Piano Sextet in B flat major

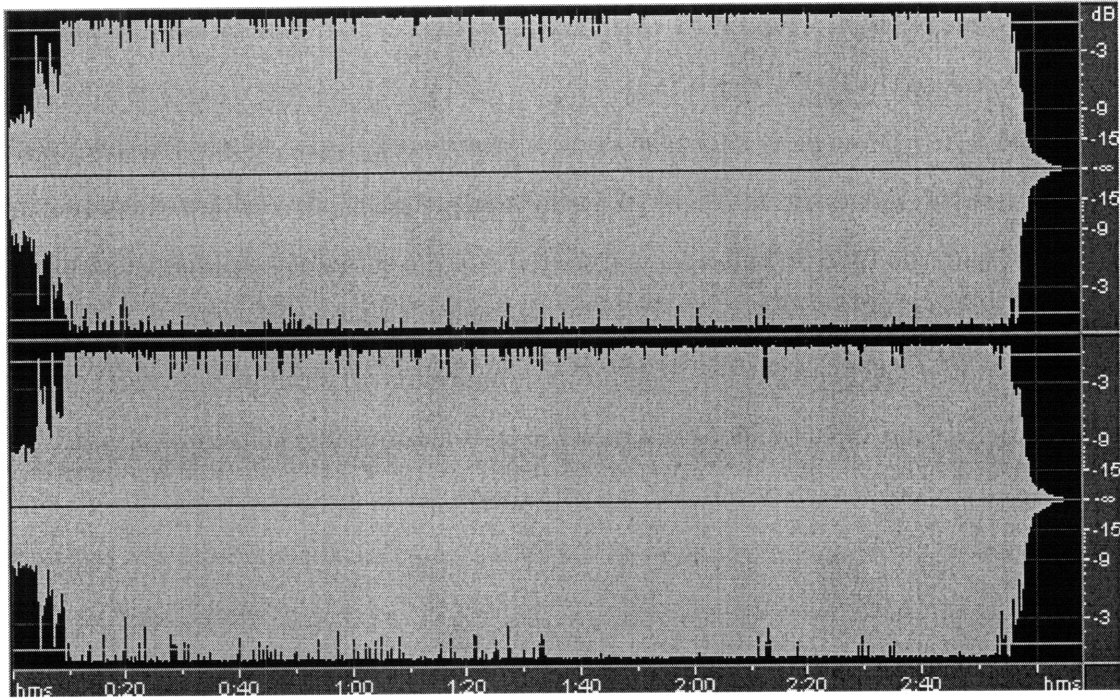


Figure 3-3: Theme Song from *Friends*

MATLAB).

This pattern suggests a dual channel architecture, for which one channel is dedicated to capturing the peak value of a signal and the other channel is more narrowly focused on the average value that composes the majority of the signal. In practice, this could be accomplished by using two identical converters, and inserting a gain stage (that multiplies a signal by an appropriate power of two) on the average value channel. Designating the average value channel 'A' and the peak value channel 'B', both channels would convert data and send the data to a digital logic device, such as an FPGA.

In the digital logic device, the two channels would be recombined according to the following rule: Channel A would be used unless an overflow condition was detected, at which time the data from Channel B would be used. Because Channel A was subjected to a gain of 2^N in order to align the converter full scale with the average signal level, the Channel A data needs to be bit-shifted to the right by N to match the values from Channel B. Thus, this technique should provide N extra bits of resolution when the signal is within its average value range ($\frac{1}{N} * \text{Full Scale}$), and

normal resolution from Channel B when the signal exceeds this threshold. Figure 3-4 illustrates this setup for $N = 1$.

Unknown to the author, this scheme was proposed by J.R. Milch [9], specifically to make improvements in the field of visual digitization. An enhancement to this technique that the author had also considered was disclosed in a patent awarded to the Eastman Kodak Company [1]. Otherwise, this technique was not mentioned in the literature while investigating resolution enhancement techniques, even though it seems that it could offer benefits in other fields, including digital audio.

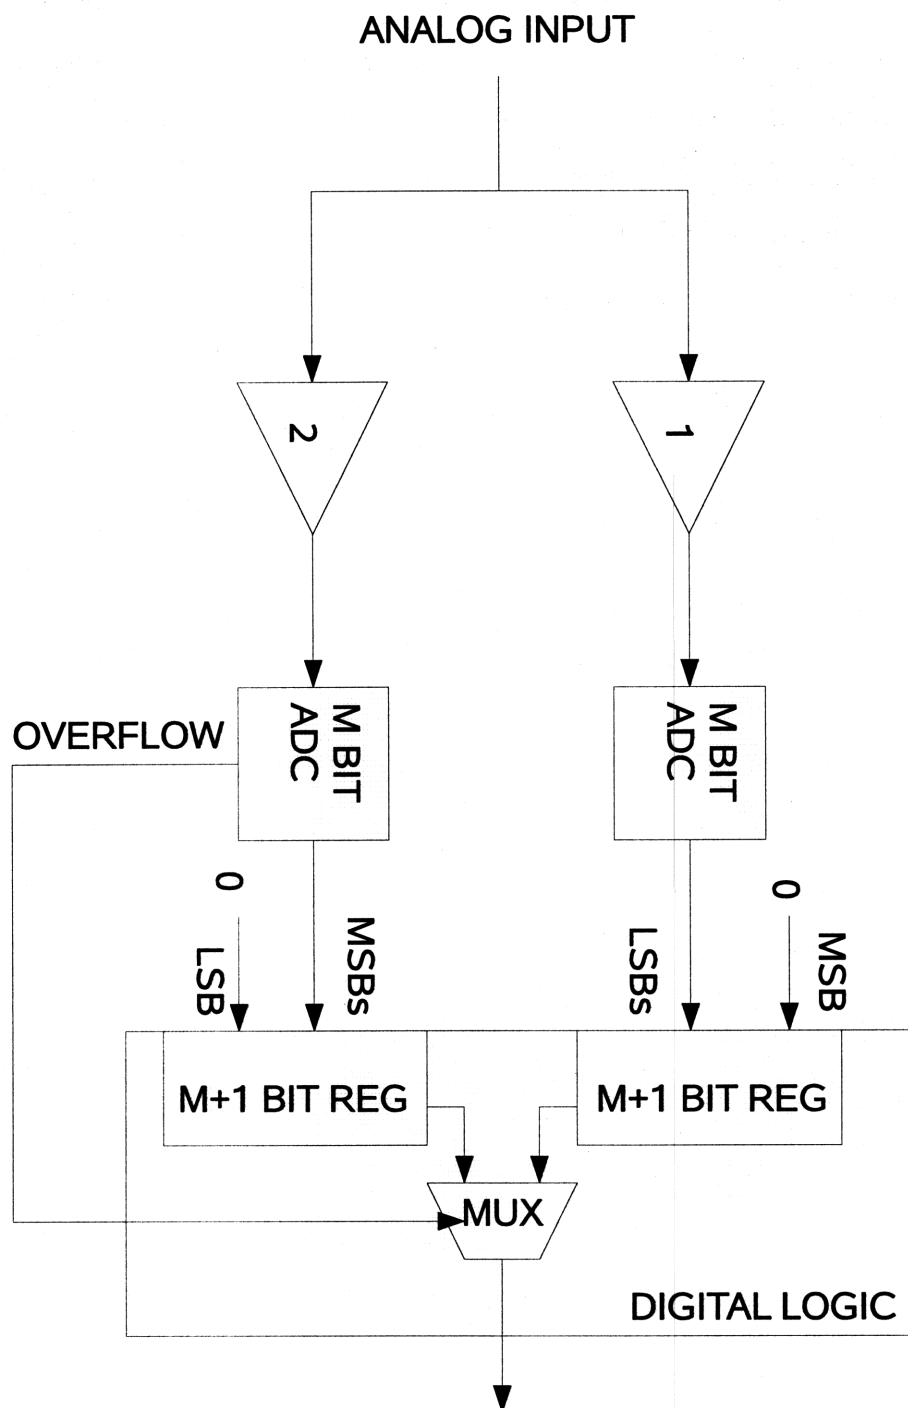


Figure 3-4: The Dual-Range A/D Conversion System

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Implementation

The Dual-Range Analog-to-Digital conversion system was realized on a proprietary board at Vtech Engineering Corporation. A diagram of the setup is provided. The A/D converters used in the system were two of the Analog Devices AD7984, a successive approximation A/D with 18 bits of resolution and a maximum sampling rate of $1.33MSPS$. For a $10kHz$ signal, the converter has a SINAD of $98dB$, indicating an ENOB of 16.¹ The primary channel had a gain stage of 2 selected. The data was stored in a SPARTAN 3 FPGA until it could be loaded into a PC.

The signal source was an Agilent 33220A Function / Arbitrary Waveform Generator. This piece of equipment turned out to be the limiting factor in analyzing converter performance, as the harmonic distortion was only specified to $-70dBc$, making accurate measurements beyond the 10 bit level impossible. For this reason, only the 8 MSBs from each converter were examined in the following analysis. The Agilent was configured to output a sine wave at $900Hz$, $1kHz$, and $1.1kHz$. These frequencies were chosen to have a period of a sine wave take approximately one thousand samples. The waveform generator was placed in burst mode with one cycle per burst, and the A/D converters began acquiring data on the trigger signal

¹While the SINAD indicates that the bottom two bits of the converter are below the noise floor, this doesn't mean they are worthless. The Spurious-Free Dynamic Range and Total Harmonic Distortion specifications both indicate performance to $110dB$, or full 18 bit performance. As previously mentioned, the noise floor can be lowered by averaging multiple samples. In the case of acoustics, the human ear is able to detect signals below the noise floor [5].

provided by the waveform generator's SYNC output. Amplitudes were chosen to overdrive Channel A to various degrees.

To analyze the converter performance, the Sine Wave Curve Fit was used to determine ENOB, as described by Lundberg [8]. 1024 points were chosen, avoiding the transients introduced at the start and end of the cycle by the waveform generator. The MATLAB Curve Fitting Tool was used to compute the Root Mean Squared Error (RMSE), by applying the Fourier fit:

$$a_0 + a_1 \cdot \cos(x \cdot \omega) + b_1 \cdot \sin(x \cdot \omega)$$

One caveat with this particular tool was that each fit calculated all four parameters: a_0 , a_1 , b_1 and ω . The parameter ω should remain one fixed frequency, however. To remedy this, the best fit ω for an undistorted sine wave should be set as the lower bound for the optimization tool to work with when fitting a clipped sine wave.

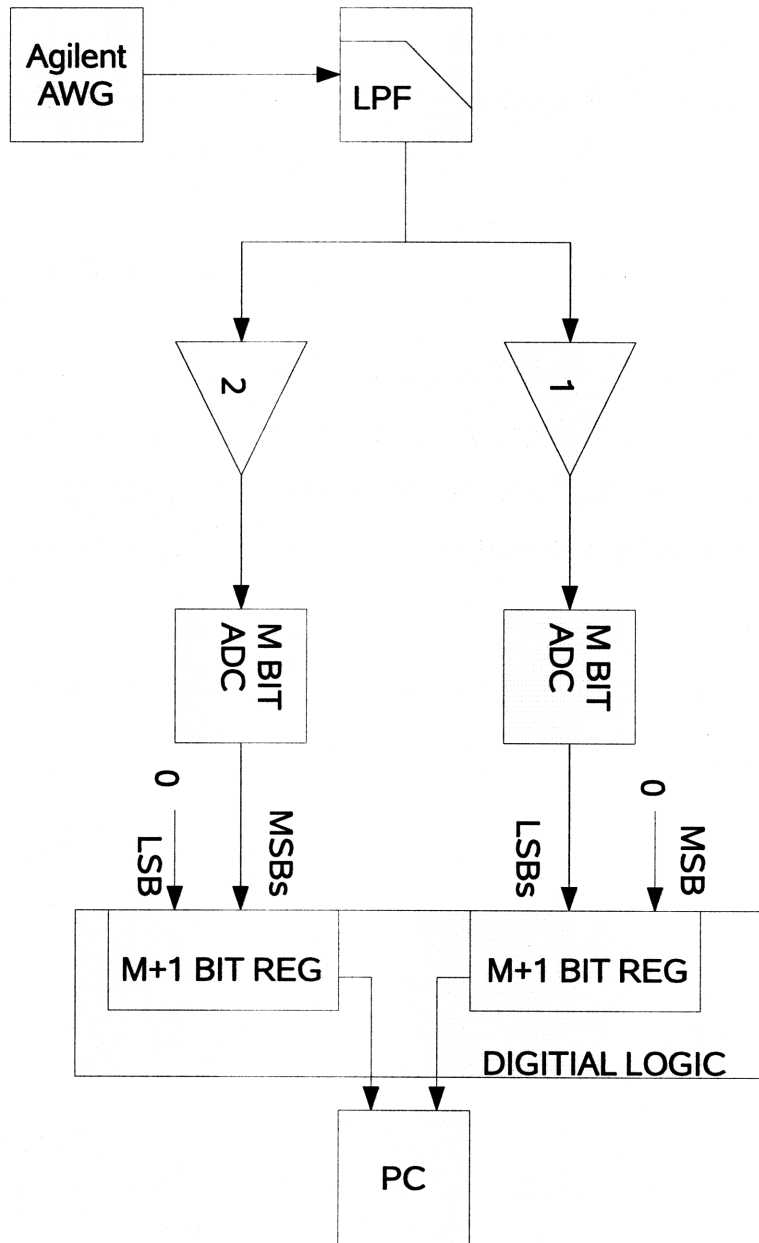


Figure 4-1: Dual Channel Conversion Experimental Setup

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Results

The performance improvements discovered at each frequency are noted, as well as the percentage of time the signal is over range. The data acquisition board had a gain common to both channels, so the output voltage of the Agilent waveform generator does not match the voltage read by the analog-to-digital converters. Also, the opamps used in the A/D converter driver circuit saturated before the voltages reached converter full scale, so clipping of the sine waves occurred earlier than expected.

Table 5.1: 900Hz Signal

Input	Samples Clipped	ENOB of Ch B	ENOB of Ch A	Combined ENOB
580mV	0	7.82	8.81	8.81
610mV	$\frac{100}{1024}$	7.81	7.67	8.49
620mV	$\frac{180}{1024}$	7.81	6.98	8.36

Table 5.2: 1kHz Signal

Input	Samples Clipped	ENOB of Ch B	ENOB of Ch A	Combined ENOB
580mV	0	7.81	8.77	8.77
610mV	$\frac{85}{1024}$	7.80	7.77	8.53
620mV	$\frac{180}{1024}$	7.83	7.08	8.38

Table 5.3: 1.1kHz Signal

Input	Samples Clipped	ENOB of Ch B	ENOB of Ch A	Combined ENOB
580mV	0	7.80	8.71	8.71
610mV	$\frac{90}{932}$	7.81	7.81	8.51
620mV	$\frac{140}{932}$	7.82	7.14	8.44

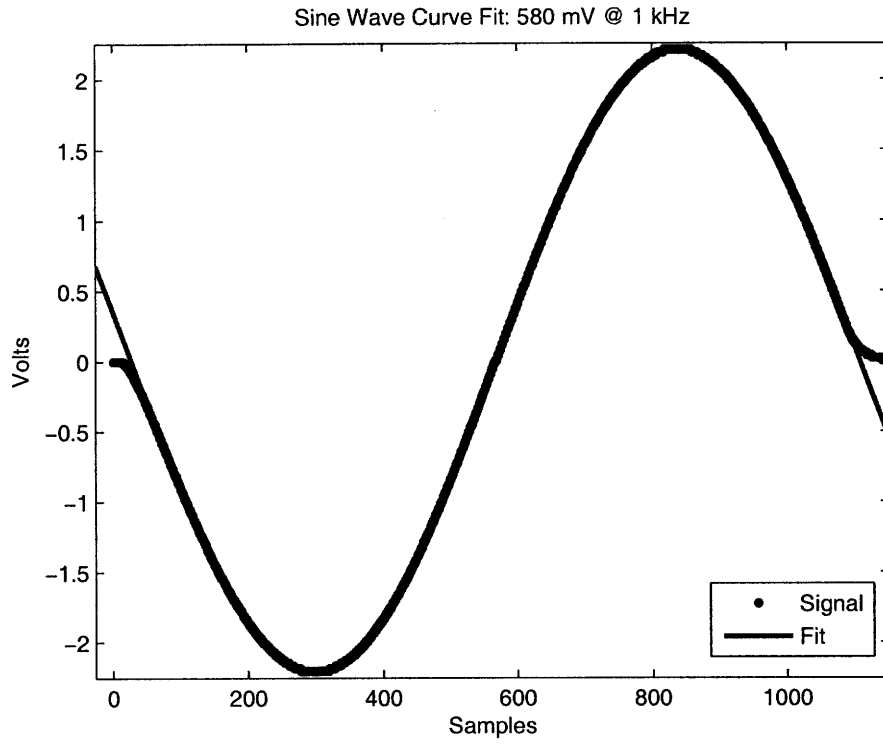


Figure 5-1: Curve Fitting: No Distortion

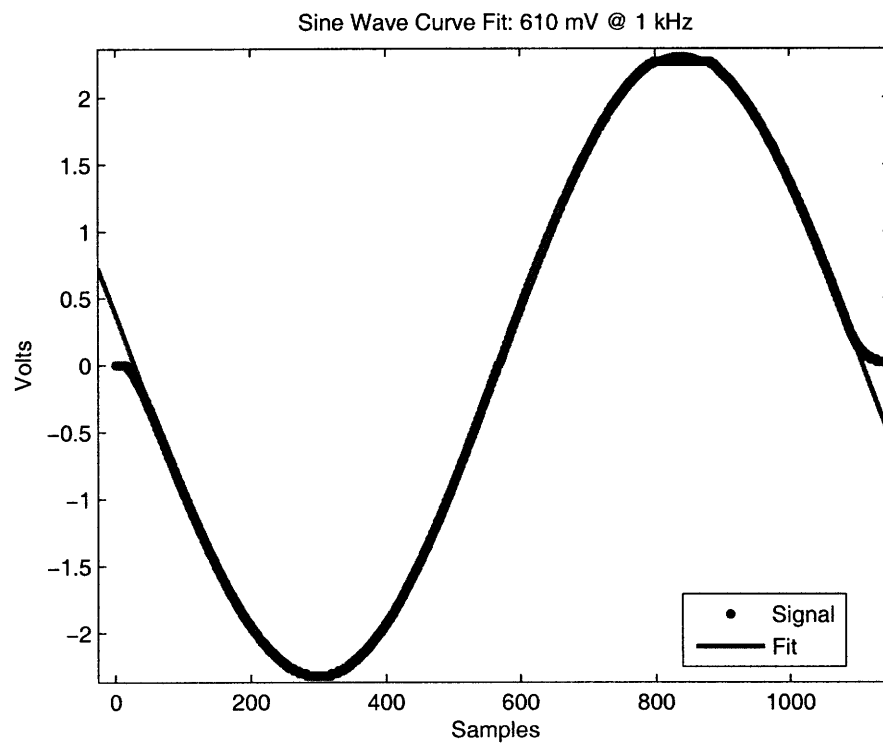


Figure 5-2: Curve Fitting: Moderate Distortion

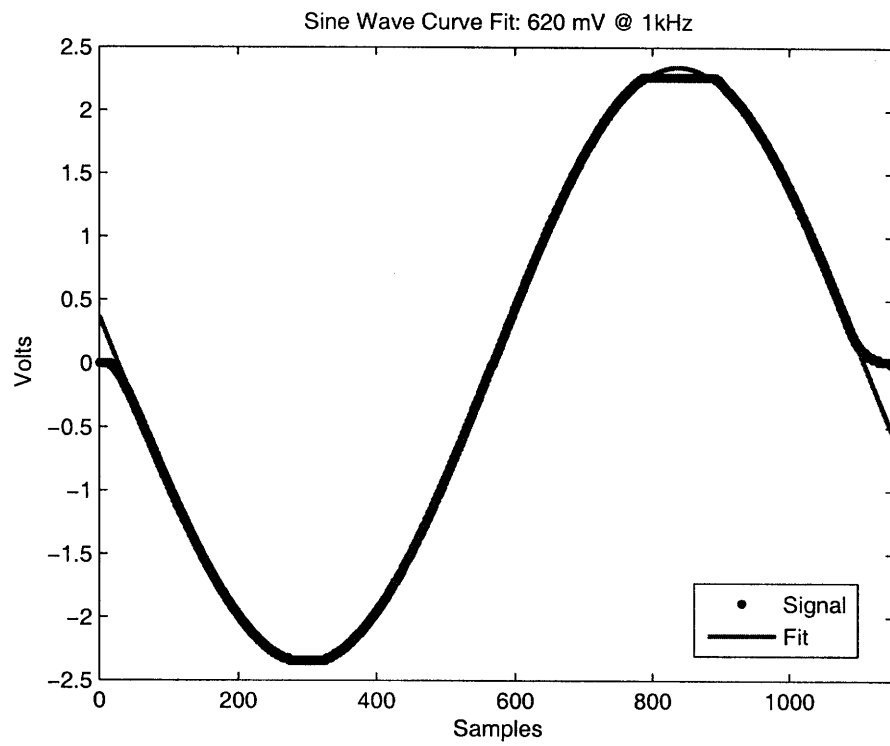


Figure 5-3: Curve Fitting: Significant Distortion

Chapter 6

Conclusion

For signals that stayed within half of the converter's range ($\pm\frac{1}{4}$ fullscale, as defined in the background section), a gain of one effective bit was seen, matching expectations. As the amplitude of the signal increased, clipping was noted on Channel A and the system performance decreased. With about 15% of the sine wave clipped, the performance gain of the combined system fell to just over half a bit (a $3dB$ improvement in SINAD). This is the same performance gain expected from time interleaving or signal averaging using two converters.

This result suggests that using a dual-range converter configuration for signals that share the characteristics of the audio waveforms examined earlier (average value within half of full scale with occasional peaking) is a valid way to increase resolution, though it is not assured to be strictly better than other methods. A system incorporating two to four converters in a signal averaging or time interleaving configuration can provide the same expected performance.

Ultimately, this technique may be incorporated into a monolithic chip design, just as the AD6645 and AD10678 are examples of monolithic chips that contain multiple A/D converters in signal averaging and time interleaving configurations, respectively.

The issue of channel matching is an important one that has not been covered in this work. This issue has been studied extensively in the case of time-interleaved converters [4], and the algorithms found should prove applicable to the dual-range

method as well.

As the patent protections on this and similar topologies will be expiring in the next three to five years, it is worth investigating the merits of this technique as an alternative to higher resolution converters that will become more and more expensive to produce as the the fundamental performance limits are approached.

Appendix A

Sample MATLAB Script

This is an example of a MATLAB script used to pull the data from the data acquisition board log file and prepare it for the MATLAB Curve Fitting Tool.

```
x620mv10khz = load('620mv10khz.log');

%loading 18 bit data into vectors
x620mv10khzchB = x620mv10khz(:,5);
x620mv10khzchA = x620mv10khz(:,6);

% Calibrating offset for both channels
ch4 = x620mv10khzchA - round(mean(x620mv10khzch4));
ch3 = x620mv10khzchB - round(mean(x620mv10khzch3));

% Calibrating gain: channel A relative to channel B
chA = round(ch4 * 2 / 2.00715);

% Truncating data to 8 bits
chBbit8 = bitshift(bitshift(ch3, -10), 10);
chAbit8 = bitshift(bitshift(ch4, -10), 10);
```

```

% Dividing channel A by 2 (Right Shift 1 bit)
chAbit8ds = bitshift(chAbit8, -1);

% Converting values to voltages
ch3volts = ch3bit8 * 10 / 2^18;
ch4voltsds = ch4bit8ds * 10 / 2^18;

% If clipping is present, it is at the extremes
a = max(chAbit8ds);
b = min(chAbit8ds);
k = 2698;

% Use Channel A unless signal is clipped
combined = zeros(k,1);
for m = 1:k
    if (chAbit8ds(m) == a || chAbit8ds(m) == b)
        combined(m) = chBbit8(m);
    else combined(m) = ch4bit8ds(m);
    end
end
end

```

Bibliography

- [1] Lawrence J. Bernstein and Kenneth A. Parulshi. U.s. patent no. 5,250,948: High resolution enhancement for dual-range a/d conversion. Filed 1991. Patent 1993.
- [2] White Flye. Successive approximation adc block diagram. Wikimedia Commons, 2006. Licensed under Creative Commons Attribution-ShareAlike 2.5.
- [3] Michael J. Guidry. A 16-bit, 100 ks/s self-calibrating cyclic analog-to-digital converter. Master's thesis, Massachusetts Institute of Technology, 2005.
- [4] Pieter Harpe, Hans Hegt, and Arthur van Roermund. Analog calibration of channel mismatches in time-interleaved adcs. *Proceedings of the European Conference on Circuit Theory and Design*, 2007.
- [5] Robert Katz. *Mastering Audio: The Art and the Science*. Focal Press, 2002.
- [6] Walt Kester. *Mixed-Signal and DSP Design Techniques*. Newnes, 2003.
- [7] Walt Kester. *The Data Conversion Handbook*. Newnes, 2005.
- [8] Kent H. Lundberg. Analog-to-digital converter testing. 2002.
- [9] J. R. Milch. High resolution digitization of photographic images with an area charge coupled device (ccd) imager. *SPIE*, 1986.
- [10] Suhas Sreedhar. The future of music. *IEEE Spectrum*, 2007.