

DIGITAL GENERATION OF LOW FREQUENCY, LOW DISTORTION TEST WAVEFORMS

by

LINLEY ELTON WOELK

B. S., Kansas State University, 1982

A MASTER'S THESIS

submitted in partial fulfillment of the

requirements for the degree

MASTER OF SCIENCE

Department of Electrical and Computer Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

1985

Approved by:


Major Professor

CONTENTS

111202 985964

List of Figures iv

List of Tables v

Introduction 1

Overview of a Digital Waveform Generator 3

 Digital Data Source and Control 3

 Analog Waveform Reconstruction 4

Sampling Theory and Waveform Reconstruction 5

 Sampling Theory 5

 Modeling a DAC as a Zero-Order Hold 6

 Waveform Reconstruction Using a Zero-Order Hold 7

 First Order Error Sources and Figures of Merit 8

 Quantization Error 10

 Harmonic Distortion 13

 Optimum Number of Samples 15

Analog Waveform Reconstruction 17

 Overview 17

 The DAC 18

 Deglitching 27

 The Filter 44

 Components and Noise 48

 The Complete Analog Circuit 56

Digital Data Source and Control 59

 ROM and Discrete Logic 59

 RAM Implementation 63

 Computer Control 64

Layout and Construction Techniques for Low Noise Systems 65

 Ground Management and Decoupling 65

 Avoiding Interference Noise 69

References 71

Appendix A: Circuit Diagram and Layout Details

LIST OF FIGURES

1. Block Diagram of a Digital Waveform Generator	3
2. A Sampled Signal in the Time and Frequency Domains	5
3. Block Diagram of a Zero-Order Hold	6
4. Impulse Response and Transfer Function of a Zero-Order Hold	6
5. Waveform Reconstruction in the Time Domain	7
6. Waveform Reconstruction in the Frequency Domain	7
7. Quantizer Characteristics and Quantization Error	10
8. Quantization Error Probability Density Function	11
9. Optimum Number of Samples per Cycle	16
10. Block Diagram of the Analog Section	18
11. Block Diagram of the DAC1146	19
12. Transfer Characteristic of a DAC	20
13. Total Quantization Error Probability Density Function	22
14. Settling Time Characteristics	26
15. DAC Glitching	28
16. A Train of Glitches and its Power Spectrum	29
17. Four Common S/H Configurations	29
18. Typical S/H Waveform	30
19. Acquisition Time	31
20. Output Waveform of a Deglitcher	32
21. Error Signals and Their Spectra	34
22. Charge Transfer	36
23. Charge Cancellation Using MOSFETs	37
24. Charge Cancellation Using an Equivalent Capacitance	37
25. Pedestal of the Integrating Configuration	38
26. Currents Causing Hold-Mode Droop	38

27. Feedthrough	40
28. A Deglitcher in the DAC Loop	43
29. Magnitude Response of a Low Pass Filter	44
30. Sallen-Key Second-Order LPF Section	47
31. Amplifier Noise Model	50
32. Typical Amplifier Current or Voltage Noise Density Spectrum . .	50
33. The Four-Pole Filter	53
34. Noise Model of the Filter Amplifiers	54
35. The Open-Loop Gain and PSRR of the OP-17 and OP-27	55
36. The Complete Analog Circuit	58
37. Block Diagram of the Waveform Generator	59
38. Timing Diagram for the Control Signals	60
39. Noise Generated by Common Impedances	67
40. Ground Loops as Noise Sources	68
41. Capacitive Coupling	69

LIST OF TABLES

1. SNR for Various Numbers of Quantization Bits	12
2. Sine Wave Distortion for Various Numbers of Samples per Wave . .	14
3. Estimates of Total DAC Noise	52
4. Estimates of Total S/H Noise	53
5. Noise Contributions of the Filter	54

DIGITAL GENERATION OF LOW FREQUENCY, LOW DISTORTION TEST WAVEFORMS

Introduction

This paper describes the generation of high-purity waveforms by direct digital synthesis. The principal motivation for such a project comes from the requirements for dynamic testing of high-resolution analog-to-digital converters (ADCs). Some of the more elegant test methods involve the analysis of the frequency spectrum of the ADC's sampled data. These methods require high-purity sine waves as the test waveform so that the characteristics of the ADC are not obscured.

The purity of the test waveform must be equivalent to at least one or two bits more accuracy than the accuracy of the ADC. For example, dynamic testing of a 15-bit ADC requires a test signal accurate to 16 or 17 bits. This is equivalent to a signal-to-noise ratio (SNR) of over 100 dB. Commercially available frequency synthesizers or function generators do not provide such accuracy. Typically, the better models specify that harmonically related and spurious signals will be around 60 dB and 90 dB below the fundamental component respectively. These sources of distortion and noise could possibly be reduced by filtering, but a more serious problem is phase noise. Phase noise appears as side bands around the center frequency in the frequency spectrum of the signal.¹ This would be much harder to reduce to the required level.

A solution to the problem is to build a special purpose signal generator that is designed for low noise applications in the specific frequency range of the ADC. The usual analog methods or phase-lock-loop methods of frequency synthesis are very complex, especially when low noise is the major criteria. A digital waveform generator, on the other hand, is relatively easy to design

and construct. In addition, this method is inherently well suited for low frequencies (less than 100 kHz) which is the relevant range for high resolution ADCs. In fact, a high-precision audio-frequency phase calibration standard which uses digital generation of sine waves was recently designed at the National Bureau of Standards.²² Another advantage is the versatility of a digital waveform generator. All the characteristics of the signal, including the type of waveform, can be digitally controlled which makes it especially attractive for automatic testing. The frequency can easily be swept through any given range within the capability of the system with no phase discontinuities and essentially zero delay time between frequencies. These capabilities make such a system suitable for many applications beyond its original purpose.

This report discusses the design of digital waveform generators. In particular, the design of a generator with accuracy to 16 or 17 bits intended to test a 15-bit ADC is given. The relevant range of frequencies, for which the ADC was designed, is from dc to about 100 Hz. First, a general overview of a system is given. Then sampling theory and waveform reconstruction are discussed. Next come topics related to the design of the system and particular components used with special emphasis on noise considerations. The digital (data source and control) and analog (waveform reconstruction) portions of the system are treated separately. Finally, some topics related to noise control in the layout and construction of the circuit are discussed.

Overview of a Digital Waveform Generator

A digital waveform generator can conveniently be divided into two principal functional blocks (Figure 1): 1) a digital section containing the data source and control, and 2) an analog section containing the elements necessary to reconstruct a signal of sufficient quality. The data source provides samples of the desired signal in the form of binary words to the analog section. Waveform reconstruction is accomplished by a digital-to-analog converter (DAC). The DAC converts the digital values to an analog signal which is then smoothed and possibly attenuated to achieve the final output. The rate at which the DAC receives new samples is referred to as the sampling frequency.

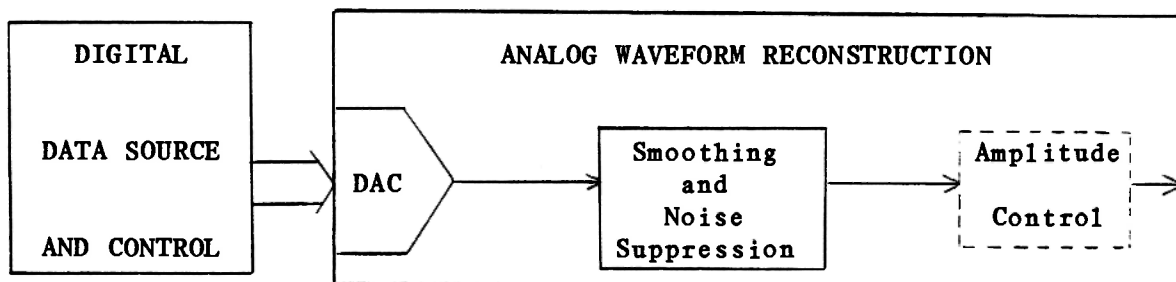


Figure 1. Block Diagram of a Digital Waveform Generator

Digital Data Source and Control

The data source generates and/or holds the samples of the waveform. It can be a computer with the appropriate interface or a block of memory with special purpose logic for control.

Memory and Control Logic - The simplest method is to use read only memory (ROM) to permanently store the values. A more versatile option is to use read/write memory (RAM - random access memory) and download the samples from a computer. In either case, one complete waveform is stored in the block of

memory. A range of frequencies is achieved using a different number of samples per wave for each desired frequency. The necessary control logic includes address generation, some method of frequency control, timing control, and data latches for the DAC.

Computer - Using a computer to generate and store the samples provides the most versatile system but requires the computer to be dedicated to the task. The digital hardware necessary is minimized with only data latches required for the DAC. An external clock may be required to insure frequency stability, in which case a first-in,first-out (FIFO) buffer may be necessary.

Analog Waveform Reconstruction

Digital-to-Analog Conversion - Reconstruction of the waveform from the samples is accomplished by the DAC. The DAC converts a digital value to an analog voltage and holds it until it is updated with a new value. In this manner it generates a staircase approximation of the signal.

Noise Suppression and Smoothing - The staircase waveform is generally smoothed by a low pass filter to achieve a more acceptable signal quality. In addition, a sample-hold is usually required to remove large spikes or glitches that are generated by the DAC and cannot be completely removed by filtering.

Amplitude Control - If desired, amplitude control can easily be included by using a four-quadrant multiplying DAC as a digitally controlled attenuator. In this application, the multiplying DAC essentially acts as a resistive voltage divider, with the division ratio controlled by the digital input.

Amplitude control in this manner is necessary because the DAC generating the waveform is usually operated with the maximum reference voltage. This minimizes the SNR due to noise sources that are independent of the signal level and may also be necessary to achieve the DAC's specified accuracy.

Sampling Theory and Waveform Reconstruction

Sampling Theory

The stored values of the signal constitute a sampled waveform. The operation of sampling is ideally modeled as the multiplication of the signal with a train of unit impulses. The sampled signal is then a train of impulses weighted by the instantaneous values of the continuous signal (Figure 2).²

The Fourier Transform of a sampled waveform is periodic in frequency with period of f_s , where f_s is the sampling frequency. Thus, the frequency spectrum of the waveform is repeated at integer multiples of the sampling frequency.²

For the continuous function $x(t)$, the sampled signal is given by $x^*(t)$.

$$x(t) = \sin(2\pi f_0 t) \quad (1)$$

$$x^*(t) = \sum_{n=-\infty}^{\infty} x(t)\delta(t - nT_s) \quad (2)$$

The Fourier Transforms of the above signals are respectively

$$X(f) = \frac{j}{2} [\delta(f - f_0) + \delta(f + f_0)] \quad (3)$$

$$X^*(f) = \frac{1}{T_s} \sum_{n=-\infty}^{\infty} X(f - nf_s) \quad (4)$$




Figure 2. A Sampled Signal in the Time and Frequency Domains

Modeling of a DAC as a Zero-Order Hold

A DAC may be modeled as a zero-order hold (zoh) (Figure 3).³ Functionally, a zoh holds the magnitude of a signal carried by an impulse function for the entire sampling period, T_s .




Figure 3. Block Diagram of a Zero-Order Hold

The zoh is essentially an integrator, it integrates the impulse function to get a unit step function. As such, it has the transfer function of a low pass filter (Figure 4).⁴ The impulse response of the zoh is given by³

$$h(t) = \begin{cases} 1 & 0 \leq t < T_s \\ 0 & \text{elsewhere} \end{cases} \quad (5)$$

The transfer function is

$$H(f) = T_s \frac{\sin(\pi f T_s)}{\pi f T_s} e^{-j\pi f T_s} \quad (6)$$




Figure 4. The Impulse Response and Transfer Function of a Zero-Order Hold

Waveform Reconstruction Using a Zero-Order Hold

When the sampled signal is passed through the zoh, the output is formed by the convolution of the input signal with the impulse response of the zoh. Modeling the sampled signal as a train of impulses, the convolution of an impulse of weight A with the unit-pulse impulse response gives a pulse of amplitude A .

$$y(t) = \int_{-\infty}^{\infty} h(\tau)x(t - \tau)d\tau = \int_0^T s A\delta(t - \tau)d\tau = Au(t) - Au(t - T_s) \quad (7)$$

The train of impulses then produces a train of pulses resulting in a staircase approximation of the sampled waveform (Figure 5).⁵




Figure 5. Waveform Reconstruction in the Time Domain

In the frequency domain, the output, $Y(f)$, is formed by the multiplication of the input, $\bar{x}^*(f)$, and the transfer function, $H(f)$ (Figure 6).⁵




Figure 6. Waveform Reconstruction in the Frequency Domain

The zoh attenuates the high frequency components. This is expected intuitively since it reduces the sharpness of the waveform in transforming impulses to step functions.

This model then does predict that the output of a DAC is a staircase approximation of the desired waveform. It provides a basis for discussing the first order sources of error in digital generation of waveforms.

First Order Error Sources and Figures of Merit

In generating the signal, the object is to get an acceptable waveform based on some criteria. The figures of merit generally used for signal quality are signal-to-noise ratio (SNR) and total harmonic distortion (thd). Thd is used to describe the effects of harmonics of the signal (components at multiples of the signal frequency) while SNR is used to describe the effects of random noise and spurious components not related to the signal. The two principal sources of error in digital waveform generation are the quantization error in the samples and the distortion caused by the higher frequency components introduced in sampling. These high frequency components are not harmonics of the fundamental signal in general but they are referred to as harmonics for convenience. The effect of the sampling harmonics is usually described in terms of thd while the effect of the quantization error is described in terms of SNR. Therefore, it is useful to have an expression relating the two.

SNR is defined as the ratio of signal power to noise power and is given in decibals (dB).

$$\text{SNR} = \frac{\text{signal power}}{\text{noise power}} = \frac{P_s}{P_n} \qquad (\text{SNR}) \text{ dB} = 10 \log \left[\frac{P_s}{P_n} \right] \qquad (8)$$

The thd of a signal may be defined as

$$\text{thd} = \sqrt{\frac{A_1^2 + A_2^2 + A_3^2 + \dots}{A_0^2}} = \sqrt{\frac{\sum_{n=1}^{\infty} A_n^2}{A_0^2}} \quad (9)$$

Thd is usually defined as some ratio that relates the amplitudes of the harmonics to the amplitude of the fundamental component but the exact definition varies with the source. Defining thd as given in (9) allows it to be related to a power ratio and hence to the SNR.

Assuming the signal is a voltage waveform passing through a one ohm resistor, the ratio under the radical represents the the ratio of the power in the harmonics to the power in the fundamental component. This is simply the reciprocal of the SNR. If the harmonics are viewed as noise signals, the signal power is the power in the fundamental component.

$$P_s = \frac{A_0^2}{2} \quad (10)$$

The noise power is the power in the harmonics.

$$P_n = \sum_{n=1}^{\infty} \frac{A_n^2}{2} \quad (11)$$

The SNR is

$$\text{SNR} = 10 \log \left[\frac{\frac{A_0^2}{2}}{\sum_{n=1}^{\infty} \frac{A_n^2}{2}} \right] = 10 \log \left[\frac{A_0^2}{\sum_{n=1}^{\infty} A_n^2} \right] \quad (12)$$

Comparing (9) and (12)

$$\text{SNR} = 10 \log \left[\frac{1}{\text{thd}^2} \right] = -20 \log (\text{thd}) \quad (13)$$

The SNR and thd thus contain the same information although in different forms.

Quantization Error

The effect of quantization error is a fundamental concept in discrete time systems.^{2 3 6} When a continuous time signal is converted to a digital quantity (quantized), it is rounded to the nearest discrete level imposed by the finite number of digits available (Figure 7). The error in this result, the difference between the digital value and the actual value, is the quantization error and has a maximum value of $\pm 1/2$ the quantization step size. The quantization error is modeled as a noise source added to the ideal signal.




Figure 7. Quantizer Characteristics and Quantization Error

The development of an expression for the mean square error (the power in the quantization noise) is based on analog-to-digital conversion where the input signal is unknown and essentially random in character. When a known signal is quantized, the input signal is not random and the quantization error could be calculated explicitly. However, for a waveform generator designed to generate any number of arbitrary waveforms, or at least a number of different frequencies of sinusoidal waveforms, it is simpler to consider the input

signal as random and derive an expression for the SNR for any possible waveform.

Since the input signal is random, the error is also random. When the number of quantization intervals is large, the probability density function (pdf) of the input signal is assumed to be uniform in each interval. Therefore, the pdf of the error is also uniform in each interval. That is, the error has equal probability of being anywhere in the interval. Thus the error is a zero-mean random variable uniformly distributed between $-\frac{\Delta}{2}$ and $+\frac{\Delta}{2}$ (Figure 8), where Δ is the quantization step size corresponding to the value of one least significant bit (LSB) of the digital quantity. In addition, if each output code is centered on its interval and if successive input samples are only moderately correlated then the quantization noise is approximately white (successive samples are uncorrelated and the power spectral density is flat).




Figure 8. Quantization Error Probability Density Function

Since the quantization error is modeled as a random variable, a measure of the performance of the quantizer must be based on a statistical quantity. The mean-square error (m.s.e.), interpreted as the average power in the noise signal, is used so that an expression for the SNR can be derived. Since the mean is zero (the power in the dc component), the m.s.e. is equal to the variance, σ , (the total average power = the power in the ac component).

For a signal of amplitude A (peak-to-peak amplitude of $2A$) quantized to an n -bit binary number, the quantization step size is given by

$$\Delta = \frac{2A}{2^n} \quad (14)$$

The m.s.e. or noise power is found by integrating the square of the error with the pdf, f_e

$$\sigma^2 = \text{m.s.e.} = E\{e^2\} = \int_{-\Delta/2}^{\Delta/2} \frac{1}{\Delta} e^2 de = \frac{\Delta^2}{12} \quad (15)$$

$$\sigma^2 = \frac{\Delta^2}{12} = \frac{(2A)^2}{(12)2^{2n}} = \frac{A^2}{(3)2^{2n}} \quad (16)$$

For a sinusoid, the signal power is $\frac{A^2}{2}$ and the SNR is

$$\text{SNR} = 10 \log \left[\frac{\frac{A^2}{2}}{\frac{A^2}{(3)2^{2n}}} \right] = 10 \log \left[\frac{3}{2} 2^{2n} \right] = (6.02n + 1.76) \text{ dB} \quad (17)$$

(Usually the 1.76 dB is dropped for convenience using only 6.02n dB.) Table 1 gives the SNR for various values of n (number of quantization bits).

Table 1

SNR for Various Numbers of Quantization Bits

n	SNR (dB)
8	49.9
10	62.0
12	74.0
14	86.0
16	98.1
17	104.1
18	110.1

The quantization SNR is the basic determinant of the system's performance. For a desired level of signal quality, the appropriate number of quantization bits must be used. Then, the maximum SNR is more or less determined by the quantization error. The quantization noise power will be reduced somewhat when the bandwidth is reduced by filtering.

Harmonic Distortion

The effect of the harmonics introduced by sampling can be determined from the equations given previously. From (6), the DAC's magnitude response is

$$|H(f)| = T_s \frac{\sin(\pi f T_s)}{\pi f T_s} \quad (18)$$

From (3) and (4), the transform of an input sine wave is

$$X^*(f) = \frac{1}{T_s} \sum_{n=-\infty}^{\infty} \frac{j}{2} [\delta(f - n f_s - f_0) + \delta(f - n f_s + f_0)] \quad (19)$$

The transform of the output is then

$$|Y(f)| = |X(f)| \cdot |H(f)| \quad (20)$$

$$|Y(f)| = \frac{1}{2} \sum_{n=-\infty}^{\infty} \left[\frac{\sin \pi(n f_s + f_0) T_s}{\pi(n f_s + f_0) T_s} + \frac{\sin \pi(n f_s - f_0) T_s}{\pi(n f_s - f_0) T_s} \right] \quad (21)$$

$$= \sum_{n=0}^{\infty} \left[\frac{\sin \pi(n f_s + f_0) T_s}{\pi(n f_s + f_0) T_s} + \frac{\sin \pi(n f_s - f_0) T_s}{\pi(n f_s - f_0) T_s} \right] \quad (22)$$

When the number of samples per wave, m , is an integer, $f_s = m f_0$. Then

$$f_0 = \frac{f_s}{m} \quad \text{and the harmonics occur at} \quad f = (nm \pm 1) f_0 \quad (23)$$

The magnitude of the components occurring about nf_s are

$$|Y[(nm \pm 1)f_o]| = \frac{\sin \pi \left(nf_s \pm \frac{f_s}{m} \right) \frac{1}{f_c}}{\pi \left(nf_s \pm \frac{f_s}{m} \right) \frac{1}{f_s}} = \frac{\sin \pi \left(n \pm \frac{1}{m} \right)}{\pi \left(n \pm \frac{1}{m} \right)} \quad (24)$$

This expression can be used to calculate the normalized amplitudes of the harmonics and the resulting thd or SNR. Table 2 gives the thd and SNR for various numbers of samples per wave, m.

Table 2

Sine Wave Distortion for Various Numbers of Samples per Wave

m	thd (%)	SNR (dB)	
2^8	256	0.70	43.0
2^{10}	1 024	0.18	55.1
2^{12}	4 096	0.044	67.1
2^{14}	16 384	0.011	79.2
2^{16}	65 536	0.00275	91.2
2^{17}	131 072	0.00138	97.2
2^{18}	262 144	0.00069	103.3
2^{19}	524 288	0.00034	109.3
2^{20}	1 048 576	0.00017	115.3

The SNR rises as the number of samples per wave increases since the staircase has more steps and approximates the waveform more closely. From Tables 1 and 2, it is necessary to use 2^{n+1} samples per wave to achieve a SNR comparable to that of the quantization error using n bits. Therefore, for some values of n it is possible to achieve an acceptable waveform without filtering. All that is required is to use 2^{n+1} samples per wave. However, this is not practical for larger word sizes. For any but the slowest signals, the required sampling frequency is beyond the capability of any available high-resolution DAC.

Optimum Number of Samples per Cycle

From Table 2, it might appear at first glance that increasing the number of samples can achieve an arbitrarily low thd. However, this is not the case since the quantization noise sets a minimum level of distortion independently. Therefore, there is some optimum number of samples that will produce the best possible waveform (ignoring the filter for the moment).

The quantization error and the harmonic distortion are independent and additive. Quantization error is the error in the step height while harmonic distortion causes the signal to be a series of sharp steps rather than a smooth waveform. In the frequency domain, quantization error is approximately white noise present in a wide range of frequencies while harmonic distortion is present only at discrete frequencies. The quantization error effectively sets a maximum SNR that can not be improved. The optimum number of samples per wave, M , will be a number that achieves a somewhat lower SNR than that given by the number of quantization bits. Then when the two sources of error are added (by the root sum of the squares), the quantization error will dominate. Comparing Tables 1 and 2, this occurs when approximately 2^{n+2} samples per wave are used where n is the number of quantization bits.

This number corresponds to the point at which the limit of the resolution in the digital samples is reached. For this, M depends on the maximum slope of the generated signal. For a sine wave

$$v(t) = A \sin (2\pi f_0 t) \quad (25)$$

the maximum slope occurs at zero and is $\frac{dv}{dt}_{\max} = A2\pi f_0$ (26)

For n quantization bits and M samples per cycle, the quantization step size, Δ , and the time interval for one sample, Δ_t , are

$$\Delta = \frac{2A}{2^n} \quad \Delta_t = \frac{T_o}{M} = \frac{1}{f_o M} \quad (27)$$

M is such that the maximum slope is equal to $\frac{\Delta}{\Delta_t}$.

$$\Delta \frac{1}{\Delta_t} = \frac{2A}{2^n} f_o M = A 2^n f_o \quad (28)$$

then $M = 2^n \pi$ or $M \approx 3 \cdot 2^n$ (29)

Using more samples, m , than the optimum number, M , (Figure 9) the resolution given by the number of bits, n , is not great enough to take advantage of the the number of samples available. All of the sample values are repeated at least once. If less than M samples are used, the waveform could be improved by inserting another sample to decrease the step size to one Δ .

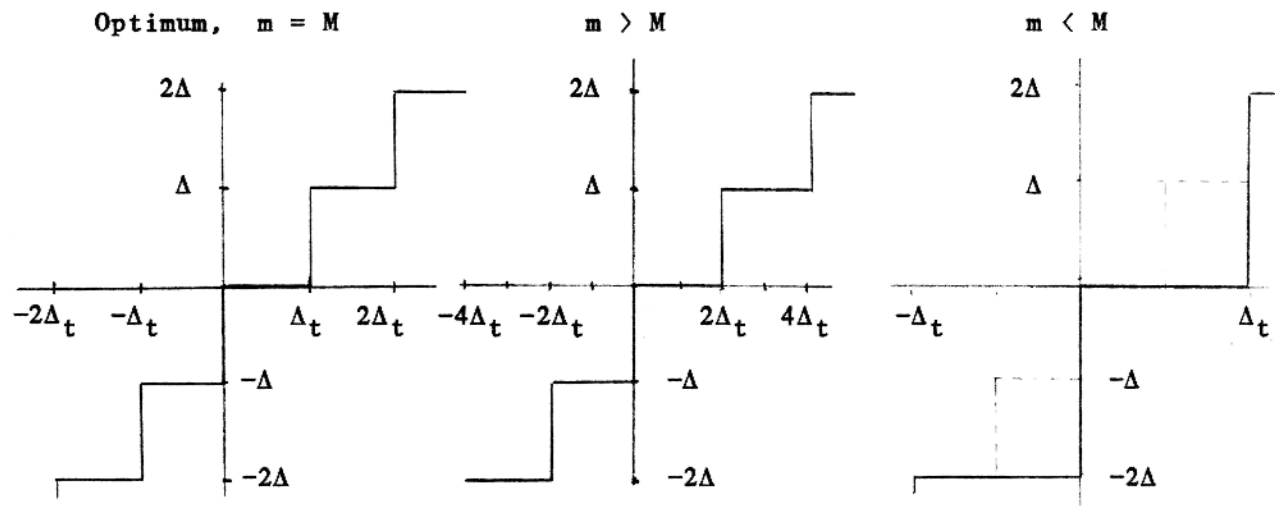


Figure 9. Optimum Number of Samples per Cycle

Analog Waveform Reconstruction

Overview

The previous discussion assumed ideal devices and the resulting first order sources of error are actually the easiest to deal with. As stated previously, the fundamental limit on the SNR of the generated waveform is determined by the quantization error after filtering. To achieve a desired SNR, all that is necessary is to use the proper number of quantization bits or word size. The harmonic distortion is more or less easily reduced by filtering, depending on the separation of the maximum signal frequency to be generated and the sampling frequency used. In this case, the filter requirements are eased considerably by the fact that the relevant range of signal frequencies is only dc to around 128 Hz while the DAC is capable of running at 50 to 100 kHz. In fact, allowing for future applications by increasing the capability to 500 Hz or even 1 kHz still leaves about 2 decades of separation in which the filter must achieve the necessary attenuation.

The characteristics of real devices, on the other hand, present the major problems. Some of the principal second order sources of error include DAC glitching, DAC linearity errors, slew and settling limitations of amplifiers, frequency stability and jitter, random noise sources in the circuit, and coupling of external noise sources into the signal. Other characteristics of real devices must also be considered, although most are less important sources of error. All of these will depend heavily on the actual devices used.

The block diagram (Figure 10) shows the principal functional blocks required for the reconstruction of the waveform. They are: 1) a latch or buffer in front of the DAC, 2) the DAC, 3) a deglitcher (a sample-hold), and 4) a low pass filter. Amplitude control in the form of a multiplying DAC

would follow the filter if it were present. These functions are discussed in this section in detail with emphasis on their error contributions.




Figure 10. Block Diagram of the Analog Section

The DAC

The function of a digital-to-analog converter is to take a digital input and convert it into an analog quantity. The DAC scales its reference voltage or current by the value of the digital input to produce the output. The output may be a voltage or current and may be unipolar or bipolar. Often the full scale range of the output may be chosen from a number of possible values. The required input code may be straight binary for unipolar output or two's complement, offset binary, true complement, or signed magnitude for bipolar output.

The DAC chosen for this system is the 18-bit Analog Devices DAC1146. This device is only guaranteed to be linear and monotonic to 16 bits, but the cost of a true 18 bit DAC is about five to six times as much. As will be shown later, the 1146 will achieve the required performance. The 1146 accepts either two's complement or offset binary for bipolar output (binary in unipolar mode). It is a hybrid module that includes a precision voltage

reference and an output amplifier (Figure 11). However, an external reference or amplifier may be used instead. The conversion is achieved by steering a current source for each bit through a resistor ladder. The current division through the ladder achieves the proper weighting for each bit. The output of the actual DAC is a current of -2 mA full scale which may be offset to achieve ± 1 mA range. The current is converted to a voltage by the output amplifier or an external op amp. Appropriate feedback resistors are included to achieve a full scale range of $+5$ V or $+10$ V for unipolar output or ± 5 V or ± 10 V for bipolar output.

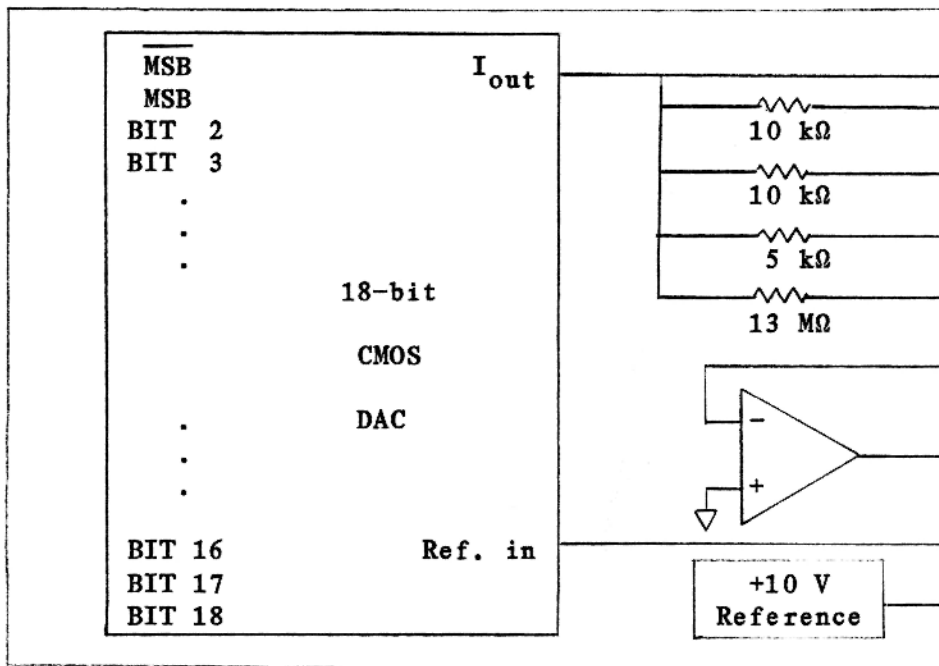


Figure 11. Block Diagram of the DAC1146 ⁸

DAC Linearity and Quantization Error - The principal measure of the accuracy of a DAC is given by the linearity specifications (differential and integral nonlinearity). These are a measure of the deviation of the output from the ideal output. As such, they give an indication of the additional quantization error introduced when the output of the DAC is not the theoretical value.

Differential nonlinearity is defined as the maximum deviation of a step

size from the ideal of one LSB (Figure 12). Differential nonlinearity is important because it gives the maximum range of analog voltages that cannot be generated. Integral nonlinearity is usually the maximum deviation from the best straight line fit to the actual transfer characteristic. It may also be measured as the maximum deviation from a straight line between the endpoints of the transfer characteristic. The integral nonlinearity represents the accumulation of the differential nonlinearity.

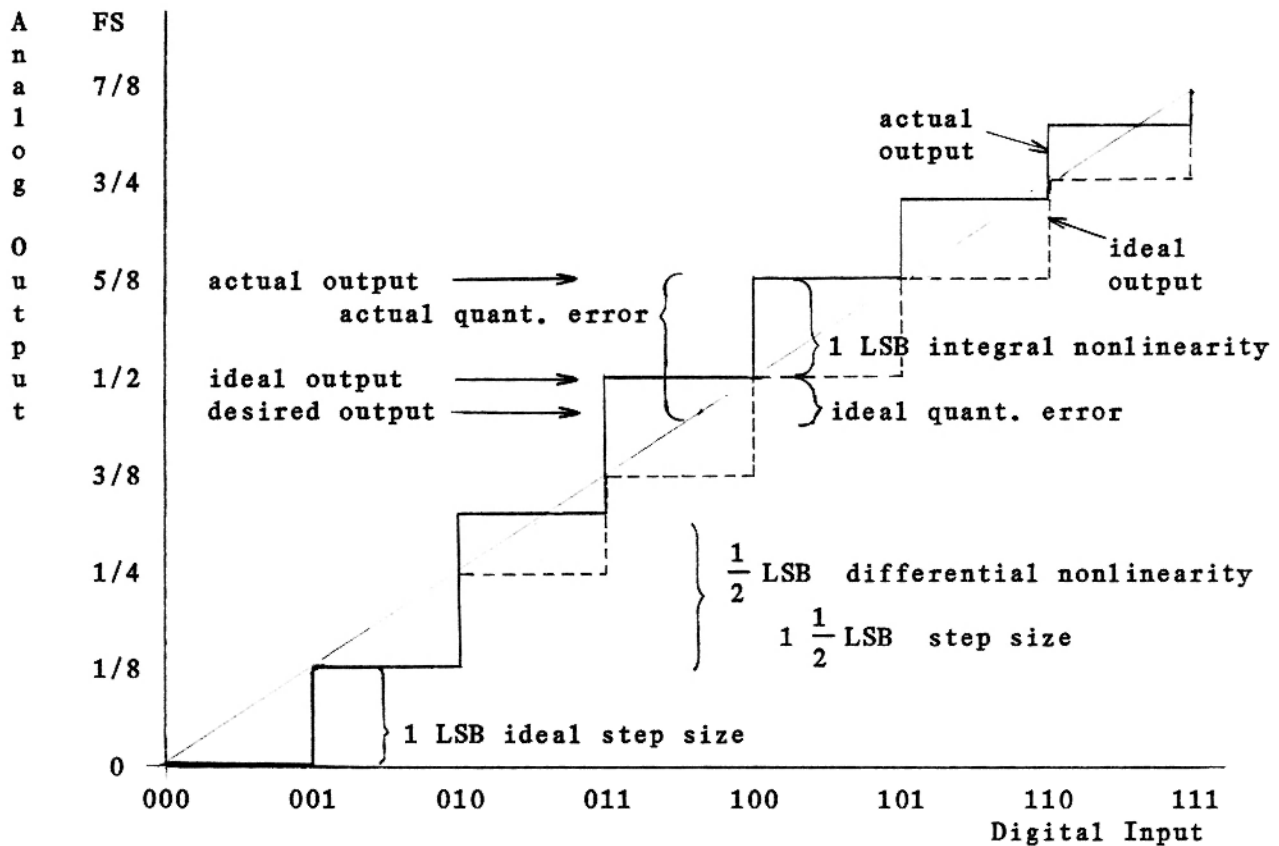


Figure 12. Transfer Characteristic of a DAC

Nonlinearity errors increase the quantization error because they cause an increase in the error between the desired output and the actual output. For example, a one LSB deviation of the output from the ideal at some input value means the quantization error could be as high as 1 1/2 LSBs (Figure 12). If the actual transfer characteristic of the DAC is not considered in the generation of the input samples, the integral nonlinearity is the true

determinant of the additional quantization error. It is easily possible for the differential nonlinearity to be a small fraction of an LSB and yet accumulate into an integral nonlinearity of one or two LSBs. The difference between the actual output and the desired output, defined as the quantization error, is clearly determined by the integral nonlinearity in this case.

However, if the transfer characteristic is known, the samples can be generated using that knowledge and the differential nonlinearity will determine the quantization error. This would require measuring the output at all possible input codes. Then the input code for a desired analog level would be chosen as the one that gives an actual output closest to this level. The maximum error would then be given by the differential nonlinearity. Whether or not this is appropriate depends on its feasibility and the requirements of the system.

In either case, the maximum quantization error is increased from the theoretical value of $1/2$ LSB or $1/2$ the quantization step size. However, it is no longer a uniformly distributed random variable. In the theoretical case, the quantization step size was uniform for each code and each code had the same maximum error. Now, the step size is not uniform and the maximum error varies for each code. The exact effect of the additional quantization error depends on whether it is determined by the integral or differential nonlinearity.

If the transfer characteristic of the DAC is not considered in generating the samples, then the integral nonlinearity determines the quantization error. The quantization noise can be viewed as the sum of two independent noise sources. The first is the theoretical quantization noise which is the random variable uniformly distributed between $\pm 1/2$ LSB. The second is the noise caused by the integral nonlinearity error. As in the case of quantizing a known signal, the integral nonlinearity is not a random quantity and could be

determined explicitly by measuring the output of the DAC at each code. But again, it is convenient to treat the noise as a random quantity. This allows deriving a single expression for the power in the quantization noise to use in evaluating the performance of the DAC.

This noise may also be modeled as a zero-mean random variable uniformly distributed between \pm its maximum value. This last assumption seems reasonable given the empirical behavior of the DAC and the large number of codes. The sum of two random variables is itself a random variable with a probability density function (pdf) formed by the convolution of the pdfs of its components. In the case of the 1146, the integral nonlinearity was found to have a maximum value of 1 1/2 LSBs. The convolution of two uniform pdfs yields a trapezoidal pdf (Figure 13) as the probability density of the total quantization noise.




Figure 13. Total Quantization Error Probability Density Function

The mean-square quantization error is now

$$\sigma^2 = \text{m.s.e.} = E\{e^2\} = \int_{-\Delta/2}^{\Delta/2} e^2 f_e de = \frac{(\sqrt{10} \Delta)^2}{12} \approx \frac{(3.16\Delta)^2}{12} \quad (30)$$

and the SNR is

$$\text{SNR} = 10 \log \left[\frac{\frac{A^2}{2}}{\frac{10A^2}{(3)^2 2^{2n}}} \right] = 10 \log \left[\frac{3}{20} 2^{2n} \right] = (6.02n - 8.24) \text{dB} \quad (31)$$

For 18 bits, the SNR is 100.1 dB. From Table 1, this is equivalent to about $16^{1/2}$ bits. The above result then shows that the nonlinearity error dominates. Since the ADC under test will have the same problem with the addition of nonlinearity error to the theoretical quantization error, its effective accuracy will be reduced to about 14 bits. The 1146 will still give better than two bits more accuracy.

Unlike the integral nonlinearity, the differential nonlinearity has very little effect on the overall quantization error. Integral nonlinearity accumulates as the input code increases or decreases in sequence and it is significant at most of the possible codes. Differential nonlinearity however, is significant only at the major transitions. Major transitions are transitions where many bits switch as the code changes in sequence. The worst transition is from 0111 1111 ... to 1000 0000 The accumulation of errors in the lower bits is usually significantly larger than the error in the MSB. In addition, the higher bits must usually be trimmed to have an opposite error that offsets the accumulated errors of the lower bits. The result is a large error in the step size at the transition. Minor transitions do not have large differential nonlinearities because the error attributed to each bit tends to be independent of the other bits (for some types of DACs). Thus, when only a few lower bits change, the error in the step height is only due to the errors in those bits. Since the error in any particular lower bit is small, the total error is also small.

The independence of the errors in each bit depends on the type of DAC.¹⁰ Some types present a variable load to the reference that depends on the code. Parallel current source DACs present a constant load to the reference since the current sources are always on. However, the currents are either steered through the ladder or through another resistor to ground. Thus, the variation in the current flowing in the ladder will result in variations in the heat

dissipated. Depending on the size of the temperature coefficients, this makes the error in each bit somewhat dependent on the code but the effect is small.

The DAC1146 is guaranteed to have a maximum differential nonlinearity of 0.00076%. This was verified by testing. More importantly, the first error of any significance was found in the eighth bit. The bottom seven bits all had an error resulting in a differential nonlinearity of less than 1/4 LSB. Therefore, any step whose transition involves only these bits will have a differential nonlinearity of less than 1/4 LSB. This is $2^{18} - (2^{18-7})$ of a possible 2^{18} steps. Alternatively, only $2^{11} = 2048$ out of 262 144 possible steps have an error larger than 1/4 LSB - less than 1%.

An expression for the power in the total quantization error could be derived by treating the differential nonlinearity as a random variable, possibly with an exponential pdf, and proceeding as before. However, a quick estimate will suffice here. Ignoring the 0.8% of errors over 1/4 LSB and simply using 3/4 LSB as the maximum of the total quantization error, the SNR for the equivalent $17 \frac{1}{2}$ bits would be approximately 107 dB. To achieve this kind of performance however, the integral nonlinearity must be included in the samples used by finding the actual output nearest the desired output.

Other DAC Characteristics - The other characteristics of DACs must be considered but turn out to be noncritical for this application - with the exception of one.

Gain errors and offset errors are not important for this application. Gain error is the deviation of the full scale output from some desired output while offset error is the deviation from zero for the zero code. Gain error changes the slope of the ideal transfer characteristic and offset error shifts the characteristic up or down. Gain error is not important because the absolute amplitude of the generated signal is not critical. Most tests are concerned with relative amplitudes. Likewise, the offset can be ignored as

long as it is known. In any case, both errors can be adjusted to zero.

Temperature coefficients are usually critical because all characteristics of a device will vary with temperature. Fortunately, for a laboratory instrument like this waveform generator, wide temperature extremes are not encountered. The self heating of the device by its own power dissipation will be the only source of temperature change. The DAC1146 uses CMOS circuits so its power dissipation is relatively low. In addition, like any lab instrument, this waveform generator should be allowed to warm up by running it for ten to fifteen minutes before it is used in any tests.

Power supply sensitivity is a dc specification given as the % change in the offset or gain per 1% change in the power supply voltage. The DAC1146 has a sensitivity of 0.001%/%. A more useful specification would be the PSRR (power supply rejection ratio) but that is not given. If the sensitivity is equivalent to the PSRR at dc, then it is very poor indeed (about 60 dB). The PSRR is usually constant from dc out to about 10 to 100 Hz where it starts to decrease so this could be a critical problem. A low frequency change in the gain would modulate the generated signal introducing a signal component at that frequency. Similarly, a low frequency change in offset would introduce a component directly into the signal. However, at low frequencies, power supplies are stable while the PSRR can be increased at high frequencies by effective filtering of the supply lines.

The settling time and slew rate of the DAC and op amp are important characteristics. They are not critical, however, since the deglitcher will determine the characteristics of the waveform presented to the filter. In general, the settling time of the DAC should be minimized to optimize the performance of the deglitcher, as will be discussed later.

Settling time is defined as the elapsed time for the output to settle within a given error band for a step input (Figure 14). The waveform consists

of three major parts: 1) an initial delay (usually very short compared to the total settling time), 2) a period of linear slewing, and 3) a final period of exponential settling to the final accuracy. The exponential settling time is the longest for high speed, high accuracy systems.




Figure 14. Settling Time Characteristics

For a current output DAC the output op amp usually dominates the settling time. The 1146, for example, settles to 0.00076% in 2 μ s for current output versus 12 μ s for voltage output for a 10 V step. The output op amp included in the module is optimized for fast settling time so it will be used. The specified time for a 600 mV step to 0.00076% is 4 μ s. Settling to 0.00076% (1/2 LSB for 16 bits) requires 7 $\frac{1}{2}$ time constants. Settling to 0.00019% (1/2 LSB for 18 bits) requires 13 time constants. Assuming the entire waveform is exponential, the settling time to 18 bits for the 600 mV step is 7 μ s. Since 600 mV will be the maximum step size, this determines the amount of time that the deglitcher will need to hold its sample. It also allows the DAC to be run at 100 kHz if desired.

A final characteristic of the DAC that is extremely important is glitching. This will be discussed in the following section with the deglitcher.

Deglitching

The second major source of error (after nonlinearity) generated by the DAC are glitches that occur at the major transitions. A deglitcher is a special sample-and-hold (read as sample-and-hold and abbreviated S/H) circuit that prevents the glitches from propagating through to the output. The S/H isolates the output of the circuit from the DAC while the DAC switches and settles to a new value. The S/H holds the previous value during this period. Using a deglitcher completely avoids the problem of glitches but introduces a whole new set of problems. The deglitcher becomes the most critical component of the circuit since it primarily determines the content of the final waveform.

Glitches - Glitches are high amplitude, low energy spikes that occur for two reasons: 1) data skew causing differences in switching times among the bits, and 2) differing rise and fall times in the internal switches of the DAC. Data skew, (bits presented to the DAC at different times) can be minimized by using latches in front of the DAC and equal length signal paths for each bit. However, this will not entirely cure the problem.

Glitches are largest at the major transitions. As an example, consider the transition from 0111 1111 ... to 1000 000 ... , a one LSB step. If the MSB arrives at the DAC last, or if the turn-on time of the internal switches is longer than the turn-off time, then the input code will momentarily go to 0000 0000 The output will slew towards -FS until the MSB finally changes (Figure 15). Then it will return to one LSB above the original value. The result is a triangular pulse (mostly). The amplitude of the glitch varies for each transition as the code and the switch times vary. In addition, the shape of the glitch will vary as exponential, overshoot, and settling effects are involved. The amplitude could be limited by using a relatively slow

amplifier, but this would result in poor signal quality also as will be shown later.




Figure 15. DAC Glitching¹¹

Since the amplitude and shape of the glitches vary with the transition, the resulting noise signal is a complex function of the signal dynamics (code transitions). The glitches will have frequency components well into the passband of the following filter and thus cannot be removed by filtering. As an example, consider the waveform where the same code transitions occur during each cycle (Figure 16). For simplicity, the glitches can be assumed to be triangular pulses. The train of glitches is made up of a number of different pulse trains of triangular shape but various heights. In this case, each of the pulse trains has a fundamental frequency twice the signal frequency. In addition, components will be present at all harmonics of this fundamental. Although some cancellation may occur due to phase differences, the distortion will still be significant. If the generated signal is well within the passband of the filter, most of the energy of the glitches will be in the passband also. While the filter may smooth the noise somewhat, it will not remove much of the energy. The power spectrum clearly shows this.

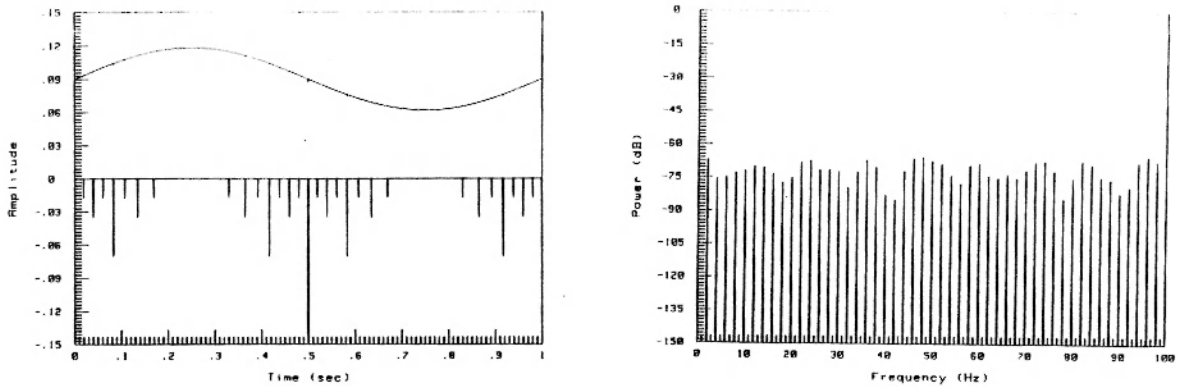


Figure 16. A Train of Glitches and its Power Spectrum

Sample-Holds - A deglitcher is a special S/H that is optimized for the particular application. The function of a S/H is to track an input signal and freeze instantaneous values when a digital command is received. Four of the many possible configurations are shown in Figure 17. Each has its advantages and disadvantages.

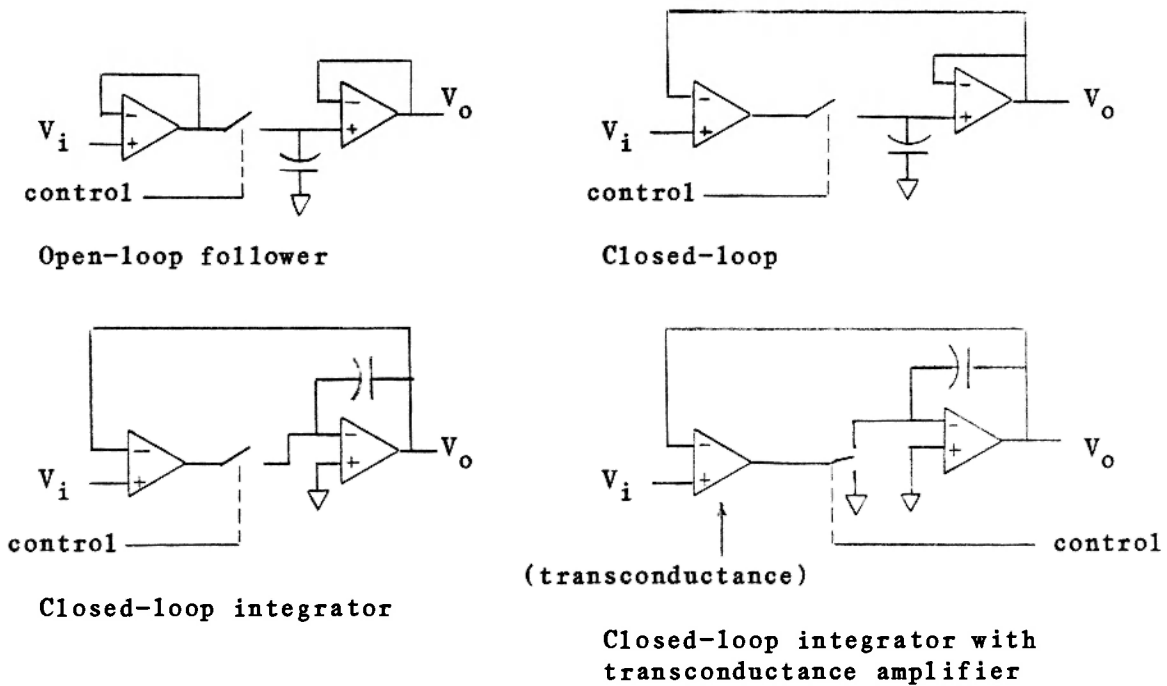


Figure 17. Four Common S/H Configurations⁹

The open loop follower is the simplest and most common type used. It provides a good example of the basic S/H. When the switch is closed, the capacitor voltage follows the input. When the switch is opened, the capacitor holds the instantaneous value of the input signal (Figure 18). The switch is usually a JFET (Junction Field Effect Transistor) or a CMOS switch (Complementary Metal Oxide Semiconductor). The amplifiers isolate or buffer the hold capacitor from the source and the load.




Figure 18. Typical S/H Waveform

A closed-loop S/H is more accurate but is generally slower. While it is tracking the signal, it essentially acts as a single amplifier. The feedback loop increases the accuracy of the system. The switch and the input circuits of the output buffer are inside the loop so errors such as switch nonlinearities and offsets do not affect the output. The integrating configurations hold one end of the capacitor and the output of the switch at virtual ground. This reduces two of the sources of error - droop and pedestal variations.

The requirements for a deglitcher are somewhat different from a normal S/H and the requirements for waveform reconstruction are different from a normal deglitcher. A S/H in data acquisition systems is primarily concerned with holding a stable voltage for a following ADC. What happens before the conversion begins is not important, as long as an accurate sample is acquired. Offset and gain errors will be important if absolute accuracy of the system is

required. A deglitcher, normally used for display applications, is intended to suppress major transients of the DAC and must not generate any of its own. However, low level spurious components or noise may be tolerated if the display is not affected appreciably. For waveform reconstruction, the frequency content of the entire output of the deglitcher is critical. As in the DAC, offset and gain errors (dc errors) are relatively unimportant while dynamic errors such as acquisition time, charge transfer, droop, feedthrough, and linearity are extremely important.

Acquisition Time - The most important characteristic of the deglitcher is its acquisition time. Acquisition time is equivalent to settling time since the deglitcher must acquire the dc signal of the DAC. The ideal case would be to have zero acquisition time so that the deglitcher reproduces the ideal staircase waveform exactly. The spectrum of this signal has no unwanted components that cannot be removed by filtering. However, no S/H can acquire a sample instantaneously and unwanted frequency components will be introduced into the waveform.

The acquisition time is defined as the elapsed time from the control signal to final settling within a specified error band (figure 19). The initial delay after receiving the command is called aperture delay time. This delay time is not important since it is a constant delay for every period.




Figure 19. Acquisition Time⁹

The actual rise and settling of the output is determined by the slew rate and current drive of the amplifiers in the circuit and by the RC time constant of the hold capacitor and the on resistance of the switch. If the amplifiers cannot provide the full initial charging current required, then slew-rate-limited charging will occur. The capacitor is charged at a constant rate with the maximum current that can be provided. If the current drive is large enough, the entire curve will be exponential. Fast acquisition time then requires high current drive, a small time constant, and a small capacitor.

The acquisition time and the shape of the waveform determine the frequency content of the waveform. Figure 20 illustrates the output waveform of a deglitcher with nonzero acquisition time. The deviation from the ideal staircase may be considered as an additive noise signal and has the form of a train of pulses of some shape.

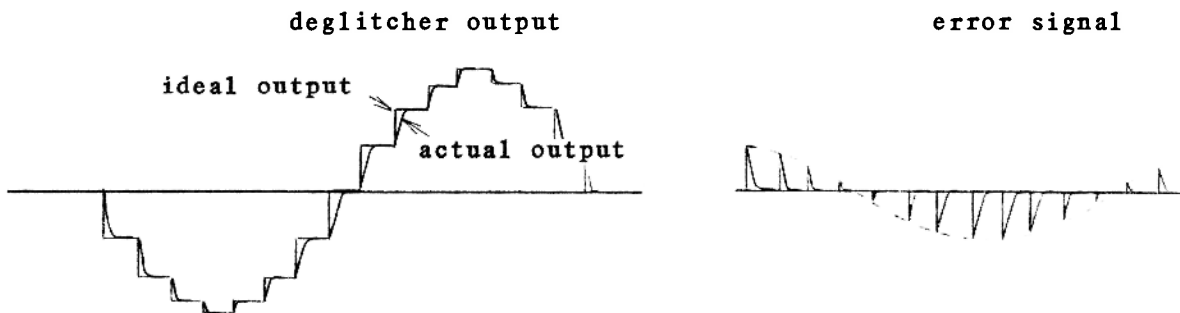


Figure 20. Output Waveform of a Deglitcher

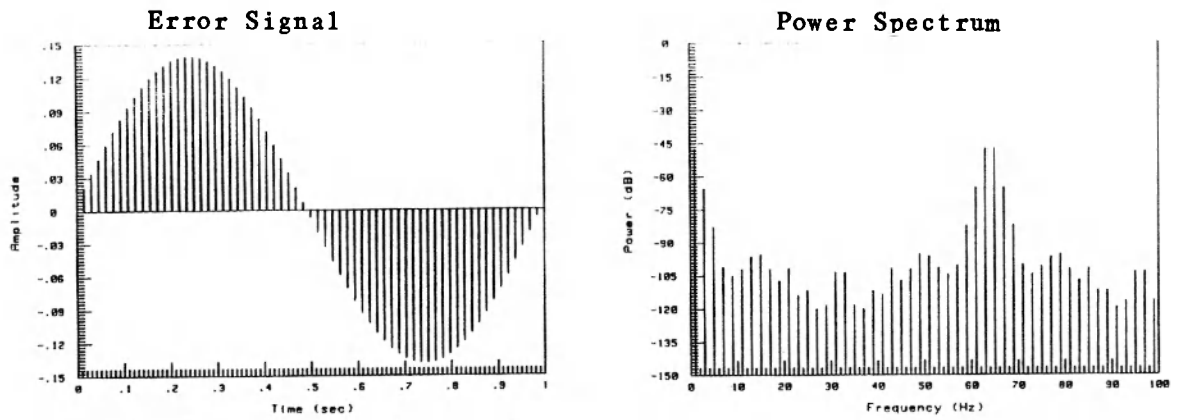
The peak amplitude of the pulses are the same as the step heights and thus are a sampled version of the signal with a phase shift and different amplitude. The entire pulses cannot be considered as a sampled version since only the peaks are related to the signal. However, if they approximate it well enough, the purity of the signal will be preserved to the extent required.

Note that it is not the shape of the pulses per se that causes the problem, rather it is the fact that only the peaks are proportional to the

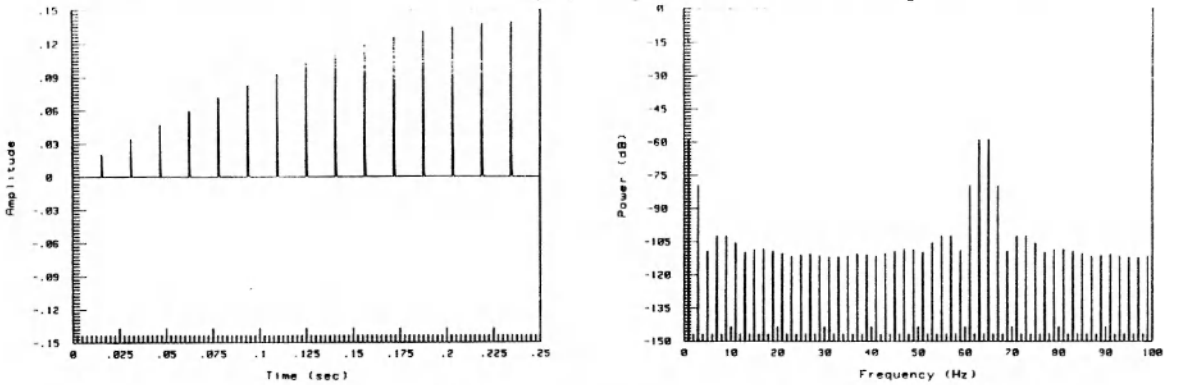
signal. Any arbitrary pulse train can be used for sampling and still allow the original waveform to be recovered exactly. The pulse train, expressed as a Fourier Series, has components only at the sampling frequency and harmonics of the sampling frequency. The operation of sampling, multiplication of the pulse train and the signal, convolves this train of impulses in frequency with the signal's spectrum just like ideal sampling. The only difference is a reduction in amplitude of the spectrum. The shape is preserved and thus the signal is preserved.

The question is then, which pulse shape best approximates a sampled signal? To answer that, the spectrum of the various shapes were analyzed -an error signal with a linear rise due to slew-limited charging and an error with a pure exponential shape. The error signals shown (Figure 21) correspond to a signal generated with 64 samples per wave. The error signals are normalized to the generated waveform in power and frequency.

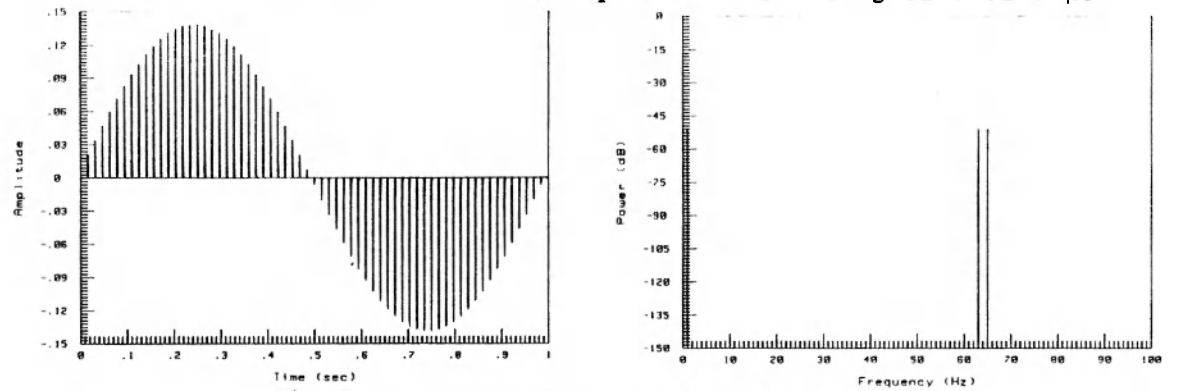
Sixty-four samples per wave would be close to the minimum used giving the largest step heights - 600 mV for the maximum step. Thus, this is the worst case that would be encountered. For a sampling frequency of 50 kHz, the sample period is 20 μ s. The first error signal (Figure 21,a), corresponds to an initial period of slew-limited charging of 2 μ s followed by exponential settling to 18 bits in 4 μ s. As the power spectrum shows, the amount of distortion this would produce is clearly intolerable. The next error signal (Figure 21,b) corresponds to an initial period of slew-limited charging of 500 ns followed by exponential settling time to 18 bits in 2 μ s. The power spectrum of this noise signal shows that it still would not be acceptable. The third signal (Figure 21,c) corresponds to a pure exponential rise time of 4 μ s. Its power spectrum is that of a sampled sine wave. Even exponential settling that requires the full 20 μ s produces no distortion (Figure 21,d).




a) Linear rise time of 2 μs, exponential settling of 3 μs



b) Linear rise time of 500 ns, exponential settling time of 2 μs



c) Pure exponential settling in 4 μs



d) Pure exponential settling in 20 μs

Figure 21. Error Signals and Their Spectra

An exponential shape for the deglitcher characteristic then, does give an error signal that approximates a sampled version of the desired signal. The shape of the exponential pulses depends on the initial peak value, which is a sample of the generated signal. If slew-limited charging occurs, in contrast, the initial slope of the error signal is the same for every sample. In addition, exponential settling will take over once the required current falls below the capability of the amplifiers. Thus, the rest of the curve will be the same for all steps and the pulse shape is independent of the pulse height. It cannot be considered as some original sampling pulse modulated by the signal.

The power spectrum shows that the exponential settling will not produce distortion. It will change the amplitude and phase of the signal. After filtering, the error signal will be a sine wave with some small amplitude and a 90° phase shift with respect to the generated signal. The two signals will sum as the following trig identity

$$A \cos(\omega t + \alpha) + B \cos(\omega t + \beta) = R \cos(\omega t + \gamma) \quad (32)$$

where

$$R = \sqrt{A^2 + B^2 + 2AB \cos(\alpha - \beta)} \quad (33)$$

and

$$\gamma = \tan^{-1} \frac{A \sin \alpha + B \sin \beta}{A \cos \alpha + B \cos \beta} \quad (34)$$

Taking $\alpha = 0$ and $\beta = 90^\circ$, this simplifies to

$$R = \sqrt{A^2 + B^2} \quad \text{and} \quad \gamma = \tan^{-1} \frac{B}{A} \quad (35)$$

Since B, the error signal amplitude, is much smaller than A, the generated signal amplitude, the effect will be slight. At any rate, the amplitude and phase are not important for the signal.

Charge Transfer - While suppressing the glitch generated by the DAC, the deglitching S/H unfortunately generates a glitch of its own (Figure 22). The digital control transition that forces the switch to open is coupled through the gate-to-drain capacitance, C_{gd} , to the hold capacitor. The resulting charge transfer induces an error voltage on the capacitor. This error voltage is called the sample-to-hold offset or the pedestal.




Figure 22. Charge Transfer

The pedestal is directly proportional to the capacitance ratio, $\frac{C_{gd}}{C_h}$, and to the magnitude of the change in the gate-to-drain voltage, ΔV_{gd} . The pedestal voltage, V_p , is approximately

$$V_p \cong \Delta V_{gd} \frac{C_{gd}}{C_h} \quad \text{or} \quad V_p = \frac{Q}{C_h} \quad (36)$$

where Q is the amount of charge transferred. The amount of charge transferred can range from less than .1 pC to over 50 pC and the pedestal voltage can range from less than 1 mV to over 100 mV. From the equations, the pedestal can be reduced by increasing the hold capacitance.

In the configuration shown, V_{gd} will vary with the signal level and thus V_p will also vary with the signal. This would be tolerable if the variation was linear. In that case, the pedestal would be a train of rectangular pulses with amplitude modulated by the signal - a sampled version of the signal again. Unfortunately, C_{gd} , also varies with the signal level and the variation of V_p is not linear. The pedestal variation will not be directly proportional to the signal. The energy in the pedestal can be minimized by

minimizing the hold time to sample time ratio. This will minimize the duty cycle of the train of pulses.

There are a number of ways to reduce the pedestal, most relying on some sort of charge cancellation technique. For example, CMOS switches may be used instead of JFETs. The complementary transistors automatically provide cancellation of the charge transfer to some extent (Figure 23). The control voltage to one transistor is inverted with respect to the control to the other. Thus, approximately equal but opposite voltage spikes are coupled into the analog signal path.¹⁴




Figure 23. Charge Cancellation Using MOSFETS¹⁴

Another method (Figure 24) uses an equivalent capacitance to collect an equal charge which is then applied to the other input of the buffer.¹² Since the charge transfer depends on V_{gd} , the hold capacitor must be kept at virtual ground requiring the use of an integrating configuration for the deglitcher.




Figure 24. Charge Cancellation Using an Equivalent Capacitance¹²

These methods will reduce the error but cannot eliminate it entirely. However, if an integrating configuration is used, the hold capacitor and the

drain of the switch are held at virtual ground. Neither C_{gd} nor ΔV_{gd} will vary with the signal level. The magnitude and polarity of the pedestal are constant and independent of the signal. The error signal composed of the pedestal is then a rectangular pulse train at the sampling frequency (Figure 25) and will be entirely removed by the filter. It will contribute a small dc offset to the signal, but again, the offset is not important. Also, since the pedestal is constant, any transient associated with the sample-to-hold transition will be more or less constant and filterable.




Figure 25. Pedestal of the Integrating Configuration

Droop - Hold mode droop or droop rate is the rate of change of the output voltage while the S/H is supposed to be holding a sample. The droop may either add to or subtract from the pedestal error, depending on the relative polarities. The voltage changes because currents bleed off charge from the capacitor. There are four sources of these currents (Figure 26): 1) input bias currents of the amplifier, 2) leakage current through the switch, 3) leakage current through the capacitor, and 4) stray currents from common node connection on the board. The droop is due to the net current flowing through the hold capacitor.




Figure 26. Currents Causing Hold-Mode Droop

If I_1 is the net leakage current, then the droop is given by

$$\text{Droop rate} = \frac{dV_c}{dt} = \frac{I_1}{C_h} \quad (37)$$

The droop rate can be minimized by increasing the hold capacitor, using FET switches and FET input op amps, and by using high quality hold capacitors. Like the pedestal, the energy in the noise signal composed of the droop may be minimized by keeping the hold time to a minimum.

As with the pedestal, the droop is important only if it varies with the signal level. Again the closed loop integrating type of deglitcher has the advantage. The bias current for the amplifiers will be constant. The main source of leakage current through the switch is the gate-to-drain leakage in a FET or the drain-to-substrate leakage in a MOSFET. These leakage currents depend on the junction voltages. Since the drain is held at virtual ground in the integrating configuration, the switch leakage current will be constant. The stray currents in the surface of the board are easily blocked with guard rings in both cases. With the other three currents held constant in the integrating configuration, the net current through the capacitor will also be constant. Thus the voltage change should be constant and filterable.

Dielectric Absorption - A third error source that occurs during the hold mode is dielectric absorption. Dielectric absorption is a characteristic of the hold capacitor. When a capacitor is charged to some voltage, discharged, and then disconnected, its voltage will creep towards the initial value. This is due to insufficient discharge time. Molecules that were polarized by the initial voltage do not return to a completely random state. Dielectric absorption thus causes the capacitor voltage to creep towards the previous sample while the deglitcher is in the hold mode. This voltage creep is a natural log function of time given by ⁹

$$V_c = K(V_o - V) \ln \left[\frac{T_h + T_s}{T_s} \right] \quad (38)$$

where K is a constant related to the capacitor, V is the value of the current sample, V_o is the value of the previous sample, T_s is the sample time, and T_h is the hold time. The effects of dielectric absorption can be minimized by using a high quality capacitor and minimizing the sample time to hold time ratio. Furthermore, since the creep is directly related to the previous sample, the error signal composed of the voltage change by itself will have the same properties as the error in exponential settling of the steps. That is, it will approximate a sampled version of the signal to some extent. Coupled with the fact that the maximum creep will be measured in microvolts, the effects after filtering should be negligible.

Feedthrough - Feedthrough is the coupling of an input signal to the output through stray capacitances in a switch while the switch is open (Figure 27).




Figure 27. Feedthrough

In deglitching applications, The concern is with the feedthrough of the glitches that the S/H is trying to isolate. For low-frequency applications such as this, feedthrough of the glitches will be negligible. The stray capacitances coupling the signal to the output are generally on the order of a few picofarads and their impedance at low frequencies is very large. Feedthrough is more important for higher frequency applications. The off isolation of switches is typically around 80 dB or greater at 100 kHz and increases by 20 dB per decade below that. The filter following the deglitcher

will have a rolloff of at least 80 dB per decade above its cutoff frequency which will be around 1 kHz. Thus the minimum total attenuation will occur in this range and will be about 120 dB. Since the glitches have very little energy in any single component, the feedthrough will be negligible.

Aperture Jitter - Aperture jitter is usually a critical specification for S/H applications. Aperture jitter is the uncertainty in the time at which the sample was actually taken and is due to the uncertainty in the switch opening time. This causes an uncertainty in the amplitude of the sample if the input signal is changing. For a deglitcher, the input signal is a dc voltage and the aperture uncertainty time of the sample-to-hold transition will only cause a small frequency jitter in the hold mode error signals. However, an uncertainty in the hold-to-sample time will cause a similar problem as that encountered in normal S/H applications. The uncertainty in the step transition will in effect cause an additional quantization error. The actual effect is an error in the time of the particular step transition. However, this may be viewed as an error in the step amplitude. The error is the difference between the actual step height and the amplitude that should be used for the actual transition time.

The magnitude of this error depends on the magnitude of the change in the desired signal during the time difference between the ideal step and the actual step. The maximum change for a sine wave occurs at the zero crossing and is equal to the maximum slew rate of the signal, $2\pi Af_0$. For a maximum signal frequency of 1 kHz with an amplitude of 5 V, the maximum error is 31 $\mu\text{V}/\text{ns}$, or 4/5 LSB (38 μV for 18 bits and a 10 V_{pp} signal). This error is then a major concern for the system. A 1 ns aperture uncertainty time adds almost 1 LSB quantization error to a 1 kHz signal. A 10 ns jitter will reduce the SNR to that of a 14 bit system.

Another consideration is the stability of the clock generating the S/H

control signal. If it is not stable, it will add to the aperture jitter. In order to keep the effects of the clock negligible, the period must not vary by more than 0.1 ns. This is 0.0005% of a 20 μ s clock period.

Configuration - From the previous discussion, the integrating configuration must be used if the sources of error are not to contribute distortion to the signal. However, closed-loop S/Hs have an additional problem. Since they keep their loop open during the hold mode, they must entirely reacquaint the signal each time they return to the sample mode. In the process, they will generate spikes during this transition. If the input buffer is a transconductance amplifier (the output is a current proportional to the input voltage), this problem will not occur. So the configuration of choice for deglitcher applications is the closed-loop integrating type with an input transconductance amplifier.⁹

Transconductance amplifiers are not generally available, although some commercial S/Hs include them. S/Hs available are optimized for fast acquisition time which usually means an initial period of linear charging in the acquisition characteristic. Since the requirement for an exponential charging curve is unique to waveform reconstruction, there is no reason to design such a characteristic into a general S/H or even a commercial deglitcher. Since the initial current required is determined solely by the input voltage and the internal resistance of the switch, this is a problem.

Commercial S/Hs are also optimized for low aperture jitter since that is a major limitation in S/H applications. In contrast, discrete switches rarely even specify the aperture jitter. In that case, a rule of thumb states that aperture jitter is around 10% of the aperture delay time.⁹ Since the aperture delay time is typically 100 ns or more, the use of a discrete switch is out of the question. A 10 ns aperture jitter would add 1 LSB of quantization error at only 100 Hz. Therefore, the deglitcher can not be made out of discrete

parts. This is unfortunate since the deglitching function could be included in the DAC's loop. The simple addition of a switch in the current path and a feedback hold capacitor in an integrating configuration would accomplish this (Figure 28). A resistor may also be necessary to assure exponential charging.




Figure 28. A Deglitcher in the DAC Loop

The S/H chosen for this system is the DATEL-INTERSIL SHM-20, a high-speed monolithic Sample/Hold. It is an integrating type and does have an input transconductance amplifier. It contains an internal 100 pF MOS hold capacitor. This small of a hold capacitor coupled with the low on resistance of the switch, assures a small time constant and fast acquisition time under normal operation. Any slew-limited charging should be avoided since the largest step size anticipated is 600 mV. In any case, the SHM-20 has a specified acquisition time of 1 μ s for a 10 V step with a rise time of 100 ns and a slew rate of 45 V/ μ s. Thus, any slew-limited charging should be limited to extremely short periods, if present at all.

The aperture jitter is specified at 1 ns allowing operation at frequencies up to 1 kHz before the added quantization error becomes unacceptable. The hold mode error specs are very good: 80 nV/ μ s droop rate, 0.1 pC charge transfer or 1 mV pedestal using the internal capacitor, and feedthrough isolation of 74 dB at 100 kHz. In addition, the open loop gain at dc is over 120 dB assuring accuracy to 18 bits.

The Filter

A low pass filter (LPF) following the deglitcher smoothes the staircase waveform to produce the final output of the system. The frequency response of a LPF (Figure 29) includes the passband where signal components are allowed through, the stopband where signal components are attenuated below some arbitrary point, and the transition band in between. An ideal filter has a rolloff with infinite slope. The cutoff frequency, f_c , is the point at which the magnitude is -3 dB for simpler types of filters.

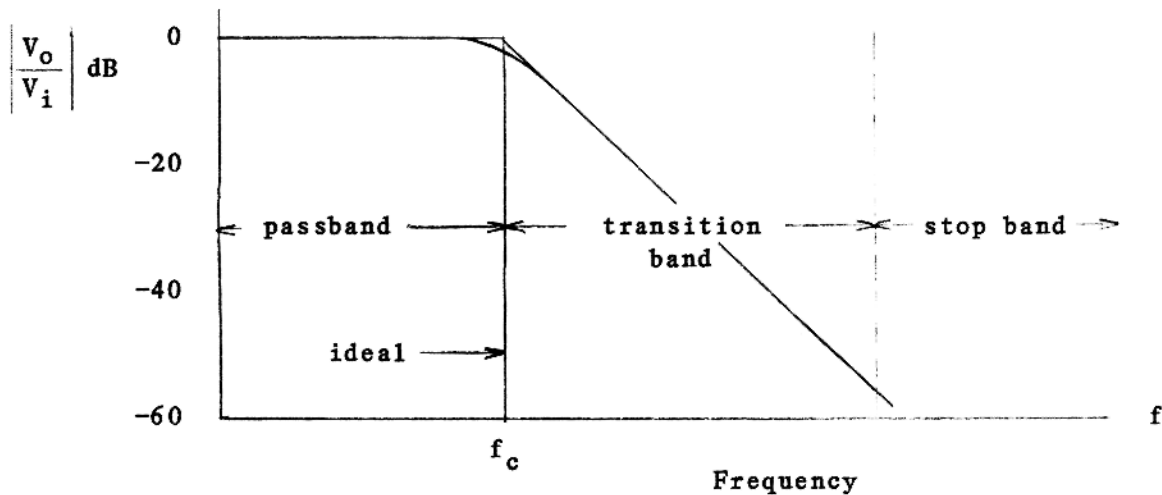


Figure 29. Magnitude Response of a Low Pass Filter

The transfer function of a filter has the general form

$$G(s) = \frac{V_o(s)}{V_i(s)} = \frac{N(s)}{D(s)} \quad (39)$$

The order of the filter is the number of poles of the transfer function and corresponds to the number of reactive elements in the circuit. In general, the slope of the rolloff in the transition band is 20 dB/decade per pole.

The phase response of the filter depends on the type of filter. The phase response is not linear for most filters - the time delay is not constant

for all frequencies. Components at different frequencies will be delayed by different amounts as they pass through the filter resulting in phase distortion. However, this is relevant only to complex waveforms with many components.

There are four standard types of filters¹⁵, each with its characteristic transfer function. All four attempt to approximate the ideal LPF in some best sense.

Butterworth - The Butterworth response has the form

$$G(s) = \frac{1}{\left[1 + \left(\frac{f}{f_c}\right)^{2n}\right]^{1/2}} \quad (40)$$

It has a maximally flat passband response since the first $2n-1$ derivatives are zero at $f = 0$. It has a more rounded cutoff than other filter types with the gain down to -3 dB at f_c . The slope of the rolloff is $-20n$ dB/decade. Butterworth filters tend to have more practical component values and less critical tolerances.

Chebyshev - Chebyshev filters achieve a sharper cutoff than the Butterworth at the expense of ripples in the passband. The rolloff slope exceeds $-20n$ dB/decade for the first few octaves but approaches that figure eventually. The initial slope can be increased if larger ripples in the passband can be tolerated.

The Chebyshev transfer function approximates the ideal filter characteristic in the sense that the maximum error is minimized. Approximating functions by polynomials is an elementary problem in numerical analysis. The polynomial that approximates a function best by minimizing the maximum error will oscillate about the function in the interval of approximation with equal amplitude ripples (the points of maximum error), thus the ripples in the Chebyshev response.

Bessel - The Bessel has a maximally flat time delay and thus achieves the minimum distortion. The phase response is very close to linear out to the cutoff frequency. However, it has a very poor magnitude response with a gentle rolloff.

Elliptic Function - The transfer function of an elliptic function LPF contains zeros as well as poles, in contrast to the other types. The zeros give infinite attenuation at discrete frequencies at the expense of return lobes in the stopband. The magnitude response has equal ripples in the passband and equal return lobes in the stopband. The elliptic function filter has the steepest rolloff of all but requires the most complex circuitry.

The steep rolloff is achieved by following a basic LPF with a high Q notch filter. The notch is set just beyond the cutoff frequency. To be effective, the notch requires a very narrow stop band, so the magnitude response soon goes back up. Thus, the first notch is usually followed by more until the original filter has achieved the necessary attenuation.

Filter Design - The design of filters is made easy by readily available handbooks. They contain component selection tables for many types of circuits and responses. The design process is little more than choosing the type of circuit and response and finding component values from the table.

The response type for this system will be the Butterworth since it is simplest and has practical advantages. There is no reason to strive for a sharp cutoff since nearly two decades are available for the transition band. The maximum frequency to be generated, determined by the limitation of the aperture jitter, will be 1 kHz and the sampling rate of the DAC will be 50 kHz. A simple 4-pole filter will achieve 135 dB attenuation at 50 kHz with the cutoff frequency set at 1 kHz. The effects of the sampling harmonics and of all the error sources with a fundamental frequency at 50 kHz will be negligible.

At the lower frequency ranges (below 1 MHz), active RC filters are used almost exclusively. Of the many possible circuit configurations, one of the simplest and best is the Sallen-Key VCVS (Voltage Controlled Voltage Source). It has a minimum number of components and is relatively insensitive to component variations. One amplifier and two reactive components produce a pair of complex poles for a second order section (Figure 30). The value of the capacitors and resistors determine the location of the poles and the response type of the filter. The frequency independent 100% feedback minimizes distortion. Higher order filters are made by cascading second order sections.




Figure 30. Sallen-Key Second Order LPF Section

Component Selection - The critical consideration here is the selection of the amplifiers. Having achieved a SNR of 100 dB so far, it would not do to introduce a distortion of 1%, a feat easily accomplished. The principal cause of distortion will be a nonlinear transfer characteristic in the op amps. The function of feedback is to reduce the dependency of characteristics such as linearity on circuit parameters. The basic relation is that feedback will decrease the open-loop nonlinear distortion by a factor of $1 \pm A_o\beta$, depending on the type of feedback.⁸ A_o is the open loop gain of the amplifier and β is the return difference or percentage of the output fed back. Thus, $A_o\beta$ is the loop gain, the difference between the open loop gain and the closed loop gain. To maintain the equivalent of 18 bits accuracy, the distortion must be kept below 0.0002 %. Open loop nonlinearity is usually less than 5 %, so the open loop gain must be above 25 000 or 88 dB in the frequency range of interest.

The op amps chosen must meet this specification.

Other op amp parameters are less important. The full power bandwidth must be larger than the desired frequency range or equivalently, the slew rate must be high enough to handle the signals. In the frequency range of interest, this is no problem. Offset and drift problems are not a concern either as discussed earlier. The other characteristic of concern is the amount of random noise generated which will be discussed in the next section.

Since the resistors and capacitors determine the poles of the transfer function, their values must be accurate or the response will not be that desired. For higher performance applications, 1% metal-film resistors and good quality 5% (or better) capacitors such as polystyrene or teflon are recommended.¹⁶

Components and Noise

There are two basic types of noise in electronic systems, device noise and induced noise. Device noise is internally generated by the components and can be minimized by proper selection of the devices. Induced noise is noise that is picked up from the outside world - primarily by capacitive and inductive coupling or through direct connections such as power supply lines. Induced noise is minimized by proper grounding, decoupling, shielding, and guarding which is discussed in the last section of this paper. Device noise is considered here.

Random Noise - Device noise is random noise generated in resistors and semiconductor junctions. The amplitude distribution is typically Gaussian (or normal) and is assumed to have a zero mean (dc offset removed). In that case, the average power (mean-square) is the variance, σ^2 , and the rms value is the standard deviation, σ . Random noise is specified by its power

spectral density since, theoretically, it has infinite power (in its infinite bandwidth). The power spectral density is the derivative of noise power with respect to frequency, in Watts per Hertz. Since power is proportional to rms voltage or current, the rms noise is also specified with units of voltage or current per (Hz)^{1/2}. The noise in a given bandwidth is determined by integrating the noise density with respect to frequency. However, useful simplifications exist that make determining noise relatively easy. Uncorrelated rms noise sources are combined by the root sum of squares - the total rms noise is the square root of the sum of the squares of the individual rms noise sources.¹⁷

There are three important types of device noise: white noise, 1/f noise, and popcorn noise. White noise is essentially constant over the frequency spectrum, thus contains constant power per hertz or bandwidth, and dominates the high frequency range. Flicker noise or 1/f noise increases as frequency decreases, contains constant power in each decade of frequency, and thus dominates in the low frequency range. Popcorn noise, the third type, is less well understood, harder to specify, and thus usually ignored. It consists of random jumping between two or more levels. White noise includes thermal or Johnson noise generated in resistances and shot or Schottky noise generated in semiconductor junctions.

Resistor Noise - The random noise generated in a resistor is modeled as a voltage source in series with the resistor. The rms value of the source is given by

$$E_n = \sqrt{4kTRB} \quad (41)$$

where k is Boltzmann's constant, T is the absolute temperature, R is the resistor value, and B is the bandwidth. At room temperature, simpler expressions may be used.¹⁷

$$E_n (\mu\text{V rms}) = 0.129 \sqrt{RxB \text{ M}(\Omega - \text{Hz})} \quad (42)$$

Since a 100 kΩ resistor generates 41 nV in a 1 Hz bandwidth, another expression is ⁸

$$E_n (\mu V) = 41 \frac{nv}{\sqrt{Hz}} \sqrt{\frac{KB}{10^5}} \quad (43)$$

Amplifier Noise - Noise in commercial amplifiers is divided into current noise and voltage noise, both including white noise and 1/f noise. The noise is usually referred to the input. The model used is an ideal noiseless amplifier with a current noise source on each input and a voltage noise source on either one of the inputs (Figure 31).¹⁷ For some devices, DACs and S/Hs for example, total noise at the output may be given.




Figure 31. Amplifier Noise Model

The current and voltage noise is best specified in a plot of noise density (Figure 32). The white noise density, N_o , and the corner frequency, f_o , can then be used to determine the noise in a given bandwidth, $f_2 - f_1$, by the following expression.¹⁸

$$N_{f_2 - f_1} = N_o \sqrt{f_o \ln \left[\frac{f_2}{f_1} \right] + (f_2 - f_1)} \quad (44)$$




Figure 32. Typical Amplifier Current or Noise Voltage Density

The $1/f$ region is primarily caused by current noise while the white noise region is dominated by resistor noise. Since current noise through resistors cause voltage noise, it is essential to keep resistances fairly small in low noise application. Also, the bandwidth should be kept to a minimum.

Once the current and voltage noise in the bandwidth is found, the total noise is determined by normal circuit analysis. The voltage noise, like offsets, is multiplied by the closed loop gain of the amplifier (noise gain). The current noise flowing through impedances produces voltage noise just as bias currents produce offset voltages. External resistor noise sources add to the output just as normal signal sources do.¹⁷

Noise Calculations - The noise specifications for the DAC and S/H are given in terms of total output noise in specific bandwidth. This is unfortunate since the given bandwidth is not useful in this application. Some approximations can be made however. The total noise is given over a wide bandwidth which is dominated by white noise. The region below 1 kHz, however, contains significant $1/f$ noise in any total. Therefore, the corner frequency must be assumed for the devices.

The bandwidth used for the noise calculations is from 0.01 Hz to 1 kHz for the DAC, S/H, and the first section of the filter. For the last section, the bandwidth is 0.01 Hz to 1.1 kHz. The upper frequency here is the equivalent noise bandwidth of the second order low pass filter found by integrating the square of the transfer function. The noise contributions of the other components pass through the entire filter. The equivalent noise bandwidth of a four pole Butterworth low pass filter is $1.03f_c$. The error in ignoring the extra 30 Hz in the upper frequency range is negligible. The lower frequency corresponds to the practical limits of testing. Any noise component with a period over 100 seconds, given its small amplitude, will essentially be dc to the device during a single test.

The DAC has a specified noise of 30 μV rms in a 100 kHz bandwidth. The bandwidth is assumed to be from 1 Hz to 100 kHz. Table 3 below gives the estimates for the noise with various assumptions for the noise corner. For a given noise corner, the white noise density is calculated using equation 44 and the 100 kHz DAC bandwidth. Then the total noise for this system, calculated using the same equation, is determined with the 0.01 Hz to 1 kHz bandwidth.

Table 3. Estimates of Total DAC Noise

Assumed f_o	Calculated N_o	DAC Noise
1 kHz	90 $\text{nV}/\sqrt{\text{Hz}}$	10 μV rms
500 Hz	92 $\text{nV}/\sqrt{\text{Hz}}$	7.5 μV rms
100 Hz	94 $\text{nV}/\sqrt{\text{Hz}}$	4.4 μV rms

The placement of the noise corner makes a large difference in the final result. A survey of amplifiers in data books will show that the noise corner for voltage noise is typically between 10 Hz and 100 Hz while the corner for current noise is usually between 100 Hz and 1 kHz. Therefore, using the estimate of 500 Hz as f_o should be a pessimistic approximation.

The S/H has a specified noise of 200 μV rms in a dc to 10 MHz bandwidth. Using equation 44 again, the white noise density, N_o , was calculated. The estimate of the noise corner has no effect on this calculation since a 10 MHz bandwidth is completely dominated by white noise. Table 4 then shows the estimates for the S/H noise in the system's bandwidth using the same three noise corners as before. Again, the figure for the noise corner at 500 Hz will be used.

Table 4. Estimates of Total S/H Noise $N_o = 63 \text{ nV/ Hz}$

Assumed f_c	S/H Noise
1 kHz	7.0 $\mu\text{V rms}$
500 Hz	5.2 $\mu\text{V rms}$
100 Hz	2.9 $\mu\text{V rms}$

The noise contributions of the filter come from the resistors in the circuit and the op amps (Figure 33). The two sets of component values illustrate impedance scaling in filters. The characteristics of the filter are not changed if all impedances in the circuit are scaled by a factor of Z - multiplying resistor values by Z and dividing capacitor values by Z (or vice versa). This allows selection of components suitable for the application.

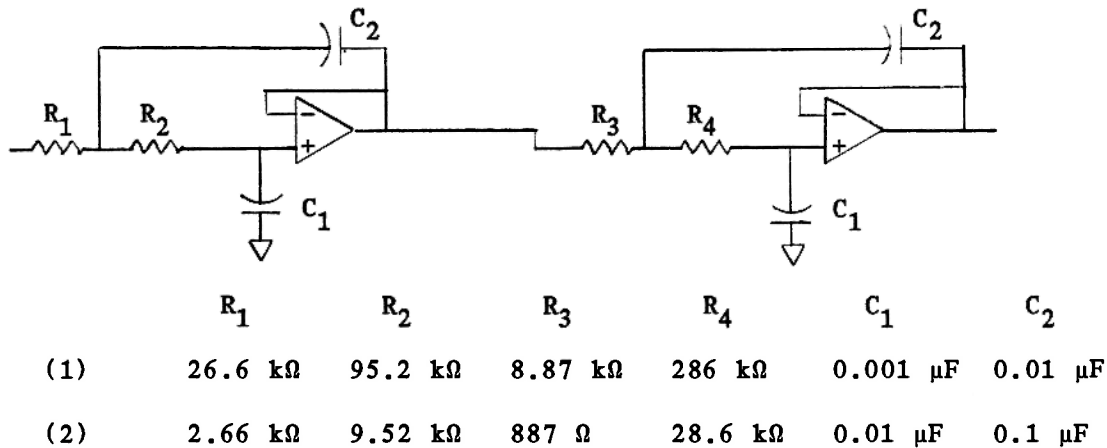


Figure 33. The Four-Pole Filter

For purposes of noise calculations, each section is essentially a unity-gain buffer. The noise model for a single section includes four noise sources (Figure 34): the current noise source on each input and the voltage noise source and resistor noise source on the noninverting input. The resistor is the sum of the two series resistors in Figure 33.

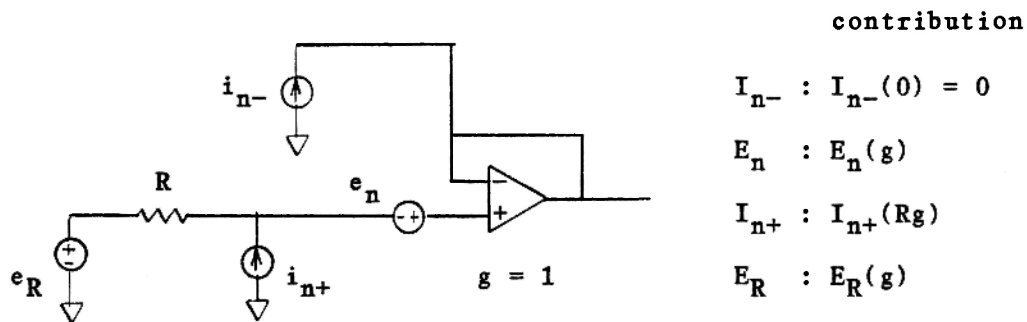


Figure 34. Noise Model of the Filter Amplifiers

Two op amps are considered, both manufactured by Precision Monolithics Incorporated (PMI). The OP-27, a low-noise precision op amp, and the OP-17, a precision JFET-input op amp. Table 5 gives the results of the noise calculations for the filter and the four cases, using equations 42 and 44.

Table 5. Noise Contributions of the Filter

Noise Source	OP-27 (1)	OP-27 (2)	OP-17 (1)	OP-17 (2)
	$\left[\begin{array}{l} e_n = 3\text{nV}/\sqrt{\text{Hz}}, f_o = 2.7 \text{ Hz} \\ i_n = .4\text{pA}/\sqrt{\text{Hz}}, f_o = 140 \text{ Hz} \end{array} \right]$		$\left[\begin{array}{l} e_n = 15\text{nV}/\sqrt{\text{Hz}}, f_o = 60 \text{ Hz} \\ i_n = .01\text{pA}/\sqrt{\text{Hz}} \text{ (negligible)} \end{array} \right]$	
	First stage: $f_2 = 1 \text{ kHz}, f_1 = 0.01 \text{ Hz}$			
I_{n+}	2.5 μV	0.25 μV	-	-
E_n	0.096 μV	0.096 μV	0.62 μV	0.62 μV
E_R	1.42 μV	0.45 μV	1.42 μV	0.45 μV
	Second stage: $f_2 = 1.1 \text{ kHz}, f_1 = 0.01 \text{ Hz}$			
I_{n+}	6.2 μV	0.62 μV	-	-
E_n	0.101 μV	0.101 μV	0.635 μV	0.635 μV
E_R	2.32 μV	0.73 μV	2.32 μV	0.73 μV
	Total Noise			
N_t	7.2 μV	1.1 μV	2.9 μV	1.2 μV

The results emphasize two points: the dominance of current noise in the low frequency range and the importance of choosing the amplifier to be used to fit the application. The smaller resistor values in case 2 gave much lower noise total since the principal source of noise is reduced - current noise flowing through resistances. The OP-27 was designed as a low noise op amp, but in this case yielded more noise because of the dominance of current noise. The OP-27 achieves low voltage noise by increasing the current in the input stage. It loses its noise advantage when high source resistance values are used. The OP-17, in contrast, achieves a lower noise total because the FET input stages make the current noise negligible.¹⁸

The OP-17s are attractive for two other reasons: high open loop gain and high PSRR (Figure 35). In contrast to the OP-27, the OP-17 will maintain the signal integrity out to 1 kHz even with an open-loop distortion of 5%. This requires 88 dB of open-loop gain above the closed-loop gain (+1 in this case). The OP-27 would start introducing distortion around 300 Hz. Similarly, the PSRR of the OP-17 will ease the requirements of the power supply filters considerably. A single-pole passive filter with a cutoff frequency of 1 kHz will counter the -20 dB/decade slope and keep the total rejection above 90 dB. In contrast, the cutoff frequency for the OP-27 would need to be around 200 Hz. Because of these advantages, the OP-17 will be used.




Figure 35. The Open-Loop Gain and PSRR of the OP-17 and OP-27²⁰

The Complete Analog Circuit

With the selection of the filter components, the analog circuit is fully determined (Figure 36) and the final error analysis can be made.

Performance Analysis - The contribution of all the error sources can be combined by converting all into power measurements and summing them to allow determining the SNR for the system. Reviewing the error sources:

- All dc errors (offsets, gain errors, etc.) are irrelevant to this application.
- DAC error sources
 - Total Quantization Error: $\text{Power} = \frac{10\Delta^2}{12} = \frac{10}{12} \frac{A^2}{2^{2n}}$
 - Temperature coefficients: Effects should be negligible if used in a lab and if allowed to warm up properly.
 - Power supply noise: Will be effectively filtered.
 - Settling time: Not important since the deglitcher determines the final waveform.
 - Glitches: Avoided by the deglitcher.
 - Random noise: 7.5 μV rms
- Deglitcher error sources
 - Acquisition characteristics: No effect since they are expected to be almost purely exponential due to the speed of the S/H and the the small step size.
 - Acquisition accuracy: Accurate to 18 bits since $A_{\text{VO}} = 120$ dB at dc
 - Charge transfer pedestal: constant pulse at the sampling frequency and therefore removed by the filter.
 - Droop: constant and too small to contribute noticeable errors anyway (.8 μV max assuming 1/2 sample period hold time).

- Dielectric Absorption: Probably no effect since directly related to signal in the same manner as the acquisition error.
 - Feedthrough: No effect. S/H isolation specified as 74 dB at 100 kHz. Assuming a typical switch isolation characteristic, it should slope up to over 110 dB at 1 kHz and below. Filter removes high frequency feedthrough.
 - Aperture Jitter: Additional quantization error of ± 0.8 LSB at 1 kHz and ± 0.1 LSB at 100 Hz. Power is $1.6\Delta^2/12$ and $.2\Delta^2/12$ assuming uniform distribution again.
 - Power supply noise: Will be effectively filtered.
 - Random noise: 5.2 μ V rms
- Filter error contributions
 - Sampling harmonics and related error signals: Attenuated below 135 dB. No effect.
 - Random noise: <1 μ V rms
 - Nonlinear distortion: Theoretically less than 0.00015 % at 1 kHz and better at lower frequencies due to high open loop gain.
 - Power supply noise: Removed by PSRR and line filters

Determining the equivalent noise power of the noise sources and summing them gives a final SNR of 99.3 dB. The quantization error clearly dominates the others. It must be stressed that the design of this system is not complete without considering external noise sources and methods of dealing with them. However, for the moment, it appears that this system can indeed give the necessary accuracy to test the 15 bit ADC. Although the final 99 dB SNR is just above that for 16 bits, the ADC will have the same problem with linearity and 16 bits should suffice.

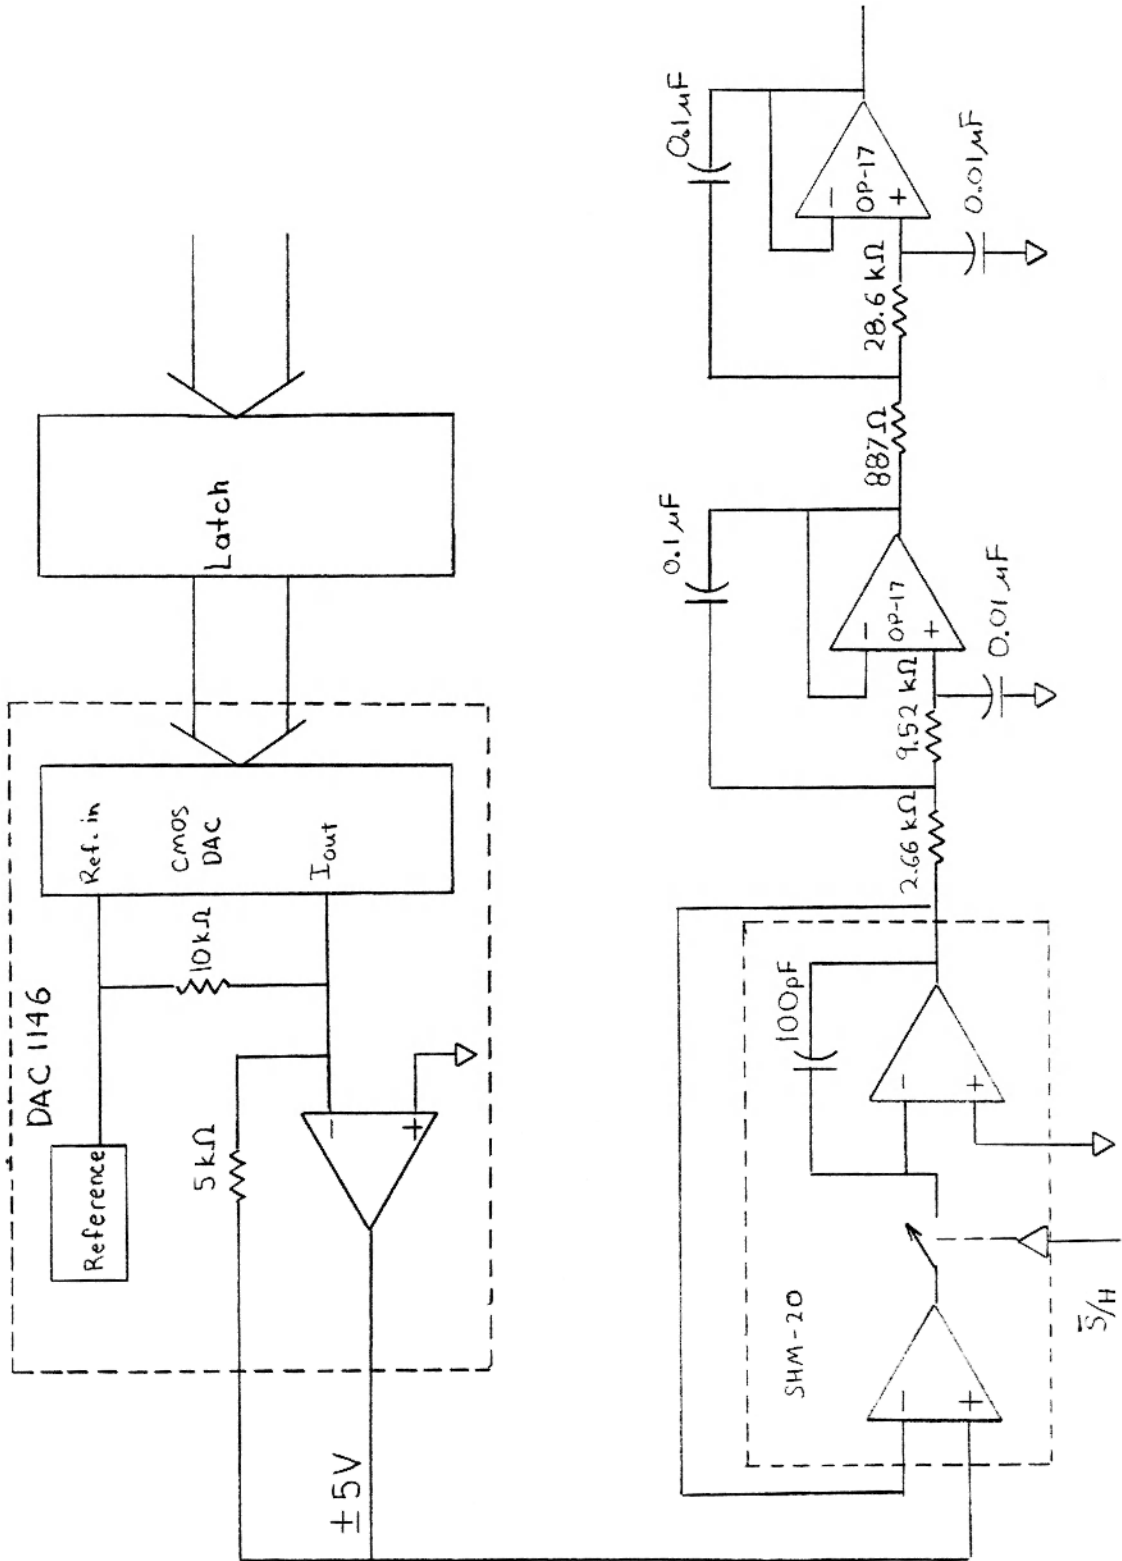


Figure 36. The Complete Analog Circuit

Digital Data Source and Control

As discussed in the overview, there are three basic options for the data source and controller: a computer, RAM and discrete logic, or ROM and discrete logic. Here, ROM is used as a generic term for EEPROM, EPROM, PROM, or ROM. Discrete logic and ROM will be used because of the simplicity of the resulting implementation. Since the waveform generator is needed immediately, and since the present requirements are for sine waves only, the more versatile configurations impose unnecessary complications. The digital portion could be modified fairly easily in the future if it is desirable. As a result, the RAM and computer options will be discussed briefly and in general terms only.

ROM and Discrete Logic

The implementation of the digital portion of the waveform generator is very straightforward. The basic scheme is to use an adder to generate the address for the ROM which contains samples of one full cycle of the desired waveform (Figure 37).

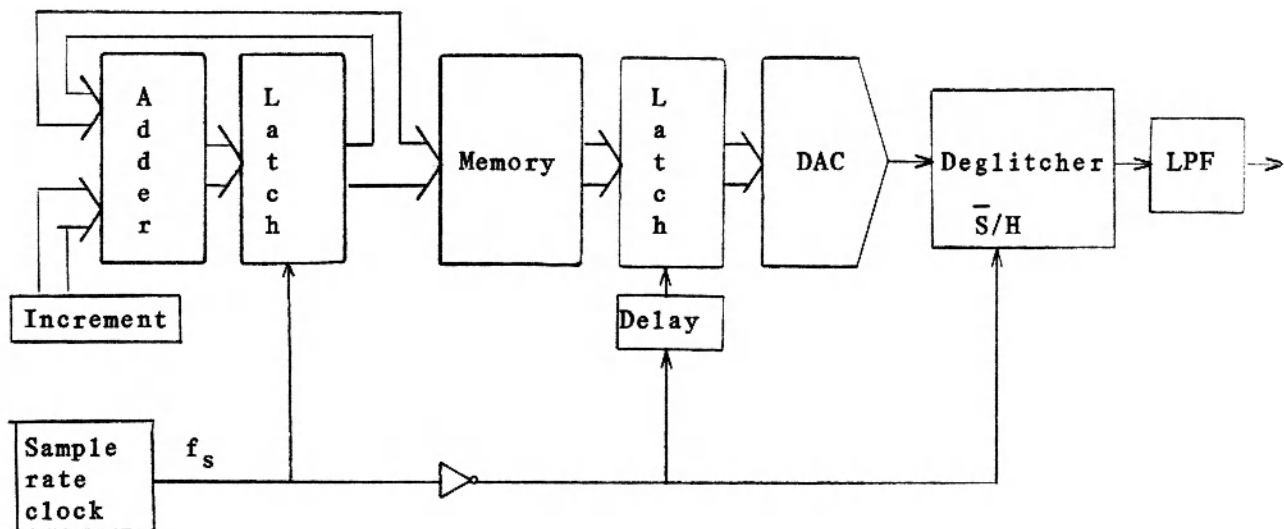


Figure 37. Block Diagram of The Waveform Generator

The adder continuously adds a fixed increment to the previous address, thus stepping through the samples in memory.²¹ When the end of memory is reached, the adder is simply allowed to overflow and the address wraps around to the proper sample in the beginning of the memory. The increment value provides the frequency control. Frequency control by this method is necessary to allow the sample rate to remain fixed and keep the sampling frequency in the stopband of the filter. Other methods of address generation, a simple counter for example, require a variable sample rate.

Timing Control - The output of the adder is latched by the sample rate clock, thus clocking out a new sample every sample period. The sample out of the memory is also latched, as required to minimize data skew in the bits received by the DAC. The sample rate clock is inverted to provide this control signal. The inverted sample rate clock also provides the control to the deglitcher. A delay of 1/2 of the sample period is encountered in the progression from the address to the DAC to the deglitcher (Figure 38). The address is generated and presented to the ROM. One half period later the sample is latched into the data latches and the deglitcher goes into hold mode. One half period later, the deglitcher returns to sample mode and acquires the new sample.




Figure 38. Timing Diagram for the Control Signals

The deglitcher must be in the hold mode before the DAC actually receives the new data, so the data latch is delayed by propagation delays. The deglitcher is in each mode for 1/2 of the sample period of 20 μ s. The settling time of the DAC should be 7 μ s maximum for the maximum step size. The remaining 3 μ s allow some margin for error. Trying to get the deglitcher out of the hold mode sooner would only add unnecessary complexity to the timing. No benefit would be gained since the hold mode errors do not degrade the signal quality. The most important reason for having all the timing signals come from a single clock signal is noise control - the fewer digital signals in the system, the better. With the timing shown, all digital transitions occur while the deglitcher is switching between modes. These have already been considered in the error analysis.

Frequency of the Generated Signal - The frequency range and resolution is determined by the number of samples contained in memory. The minimum frequency corresponds to an increment value of one - every sample in the memory is used for each cycle of the waveform. The maximum frequency will have the largest increment value, using the fewest samples per cycle.

- $f_s = 50 \text{ kHz}$ fixed sampling frequency ($T_s = 20 \mu\text{s}$), clock for the system
- $f_o =$ signal frequency
- $S =$ memory size, (number of words in memory)
- $I =$ increment value
- $N =$ number of samples per wave

The signal frequency is given by the sampling rate divided by the number of samples per cycle.

$$f_o = \frac{f_s}{N} \quad (45)$$

The number of samples per wave is given by the memory size divided by the

increment value.

$$N = \frac{S}{I} \quad (46)$$

From equations 45 and 46, the signal frequency can be expressed in terms of the increment and the memory size.

$$f_o = \frac{f_s \cdot I}{S} \quad (47)$$

The minimum frequency, occurring when I is one, is f_s/S . From equation 47, the signal frequency can only be a multiple of the minimum frequency. This minimum frequency then is also the frequency resolution of the system.

For a sample rate of 50 kHz, a minimum frequency and resolution of 1 Hz would require 50 k of words in memory. Assuming the use of Nx8 ROMS, this is 150 kbytes - a little extravagant for a simple function generator. A more realistic figure is 8 k or 16 k words (24 kbytes and 48 kbytes). Using 16 k words gives a minimum frequency and a resolution of 3.125 Hz while using 8 k words gives a minimum frequency and resolution of 6.25 Hz. The current ADC to be tested must meet specifications for a frequency range from dc to 128 Hz. The ADC can reasonably be expected to be more accurate at low frequencies, so testing at around 3 Hz to 6 Hz and at dc (static testing) should characterize it sufficiently. If lower frequencies are absolutely required, the filter cutoff frequency and the sampling rate can be lowered.

Memory Options - The amount of memory depends somewhat on the memories used. The chip count should be kept to a reasonable figure. Currently, K-State only has equipment to program 2716 or 2732 UV erasable EPROMS, both of which are readily available. Using 2716s (2k x 8) is not a reasonable solution. If 2732s (4k x 8) are used, 12 chips would be required for 16 k words and 6 chips for 8 k words. A better solution would be to use 2764s (8k x 8). Only 6 chips would be required to get 16 k words and the double memory capacity comes at only a 50% increase in cost. Using these, however, depends

on the expected acquisition of a new PROM programmer for the department. Due to the uncertainties, this system will use 2732s and 8 k words requiring 6 memory chips. It can be upgraded in the future if desired.

Generating the Samples - The generation of the sine wave samples and the subsequent programming of the EPROMs is a problem that has not been dealt with yet. Obviously, the current hand entry system via the North Star is out of the question. A program will have to be written that generates the samples and allows them to be accessed for programming the EPROM. The trick will be to get the proper bits from each sample in the three separate EPROMs. It is also important to remember that the last sample and the first sample must be in sequence, i.e. the last sample should not repeat the first.

Characterizing the DAC and using the information would improve the system significantly as stated earlier. This should be fairly easy to do on the HP test system. A good method for putting the data in simpler form would be to determine ranges of adjustment. Since the seven lowest bits have relatively small errors, there should be fairly large intervals in the transfer characteristic where the integral nonlinearity is fairly constant. Large jumps should occur only at the major transitions. For example, from 0... 0000 0000 to 0... 0111 1111, the integral nonlinearity should be below 1/2 LSB. Thus in this range, no adjustment to a sample would be needed. In another range, the sample might require an adjustment factor of -1 LSB. In this manner, the measurements would not need to be stored, only the ranges and adjustment factors would be needed.

RAM Implementation

Using RAM instead of ROM is basically an extension to the system just described. An interface to a computer system for downloading the waveform is

all the hardware that would be added. The method of choice for downloading would be through a parallel interface. The HP9845 would be the logical choice for the computer to use since it would also be running the tests. Three-state latches would be necessary to allow getting the address and all 18 bits of data from the 16 bit interface before the data was written into the memory. The latch collecting the adder output would have to be a three-state latch so the computer could access the RAM. Static RAM would be the best choice since refresh circuitry would not be required.

The principal advantage of using RAM would be the ability to generate any arbitrary waveform. This configuration would be just as flexible in that respect as using a computer to directly supply the data. Once the data was downloaded, however, the computer would be free to do other things, such as run the test and collect data from the ADC. Thus, only one computer would be required in the test system.

Computer Control

The computer control option provides the maximum versatility for the system. Besides allowing different waveforms to be generated like the RAM, the parameters of the wave could be changed at will. Another advantage is the minimization of the hardware, especially memory. The amount of memory required for the waveform generator is about as much as a typical personal computer has, so it makes sense to use the memory in a computer. In fact, the computer could generate the samples as needed (if it was fast enough) and not use memory at all. The principal disadvantage is the requirement for the second computer in the test system. Besides the inconvenience, a second noise generator next to the waveform generator could only hurt the performance of the system.

Layout and Construction Techniques for Low Noise Systems

The previous discussion of noise and error sources in the system was concerned with only the noise generated by devices in the analog subsystem. The second major concern is noise that is introduced into the analog circuits from external sources. This noise can be divided into two types: direct coupled noise and interference noise. Direct coupled noise is noise that is injected through direct connections. It is dealt with by proper ground management, decoupling, and circuit construction. Interference noise is induced into the circuit by capacitive and inductive coupling. It is reduced by effective shielding, guarding, and circuit construction techniques. External noise is intimately connected with the physical layout and construction of the system. In fact, proper layout and construction is one of the best methods for dealing with the noise.

This circuit has two intrinsic advantages in reference to external noise: 1) the bandwidth is limited to 1 kHz so the filter will do much to reduce the interference noise, and 2) the digital circuits switch at one basic frequency so the transitions occur only during two basic states of the analog circuits. The digital portion of the circuit is considered an external source when dealing with the analog section.

Ground Management and Decoupling

The two principal sources of noise discussed here are the digital circuits and the power supply leads. The digital circuits have a direct connection into the analog system through the data lines to the DAC and the control line to the deglitcher. The digital and analog portions are also connected through the power supply and ground leads. Effective control of

these noise sources is primarily concerned with decoupling and filtering the incoming signals. In addition, care must be taken not to generate additional noise sources in the system by common impedance connections. This is avoided by good ground management and power distribution techniques.

Before any of these methods are discussed, the cardinal rule of low noise systems should be noted - **USE CMOS DIGITAL COMPONENTS.** The large current spikes drawn from the power supply and returned to ground by TTL gates are a troublesome and absolutely unnecessary noise source.

Decoupling - Decoupling is the bypassing of unwanted signals in the supply lines around a component. Normally, a capacitor is placed between the supply line and the desired return point. The capacitor supplies a low impedance path for high frequency signals and thus diverts them from the terminals of the component. The action can be improved by inserting a small resistor in series with the supply line. The resistor gives a lower cutoff frequency for the filter and dissipates some of the power rather than just sending it somewhere else. Resistors must be used with care, since variations in load current will modulate the supply voltage to the components.

Decoupling and filtering not only keep unwanted noise from entering a circuit, they also keep noise generated in the component from leaving via the supply lines and bothering other components. This is especially important for digital circuits since they generate large current spikes when switching.

For the digital circuits, the standard 0.1 μ F decoupling capacitor for every couple of chips will do nicely.

The analog circuits are a different story. Good PSRR specifications are given only for the OP-17s. The DAC specifies a poor power supply sensitivity and the S/H has a specified PSRR of 65 dB for some unknown frequency - probably dc. Thus, filtering of the supply lines to these two components will be necessary. This is accomplished by two RC filters on each supply lead - two

10 Ω resistors with a 100 μF Tantalum and a .1 μF ceramic capacitor. The tantalum provides the low cutoff frequency while the ceramic capacitor is necessary for higher frequencies where the intrinsic inductance of the tantalum capacitor causes its impedance to increase.

The two poles and eventual -40 dB decade slope should serve to keep the total PSRR of the DAC and deglitcher high enough for noise suppression. The first pole occurs at about 16 Hz, so the PSRR of the devices will be increased by 20 dB at the first major noise frequency on the supply lines - the 120 Hz ripple. At the high frequencies where the digital noise will occur, the two poles should be more than adequate.

The OP-17s have an adequate PSRR out to 1 kHz. A single pole filter with a cutoff frequency set in this neighborhood will counter the -20 dB/decade slope of the PSRR and keep it at around 100 dB.

Ground Management - The placement of the filters in relation to the components is as important as their being present. The primary concern is the path taken by the noise currents on their return to the power supply. When these currents flow through lengths of ground conductor which have some nonzero impedance, noise voltages will be developed. If these lengths of conductor also form part of the return path for the signal, then the noise is effectively a source in series with the signal (Figure 39).




Figure 39. Noise Generated by Common Impedances

This noise generated by common impedances is a result of the true nature of "grounds". For noise purposes, a ground is properly defined as a low impedance path for the return of currents to their source. The ideal definition as a common reference point only serves to obscure the problem.²³

For low frequency systems, proper ground management includes separate return paths for each circuit. The paths are connected at only one point - the single point ground. Series connections of grounds should be avoided. In larger systems, the resulting wiring requirements may be unreasonable and a compromise is often made. Grounds from different groups or types of circuits are kept separate but locally a common ground path is used. For example, signal returns are kept separate from high current load returns.

The same principles must also be applied to power distribution systems. Connecting the power to components in a daisy-chain makes the voltage at the end of the chain dependent on the current being drawn from the other components. The common impedance path again creates a noise source applied to the power terminals of the components and the IR drops may reduce the voltage below rated levels.

A second major concern in ground management is creation of ground loops which contain potential differences. The potential difference again is essentially a noise source applied to the loop (Figure 40). This is much the same problem as common impedances and is avoided by the single point ground system. Ground loops also are susceptible to noise picked up in magnetic fields.

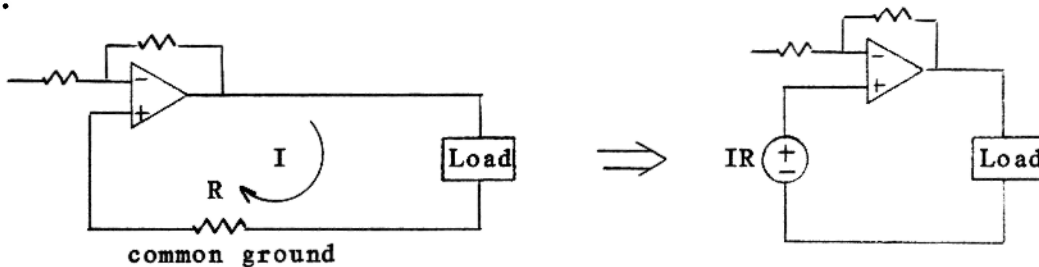


Figure 40. Ground Loops as Noise Sources

Avoiding Interference Type Noise

The two paths that allow interference noise to enter the circuit are inductive and capacitive coupling. Capacitive or electric coupling occurs through stray capacitances between conductors in the circuit. Every pair of conductors in proximity to each other forms a capacitor with some nonzero capacitance. Effective shielding can be employed to eliminate these stray capacitances. Inductive or magnetic coupling occurs whenever a loop exists in a magnetic field. Conductors carrying large currents will create significant magnetic fields within the circuit. Avoiding magnetic coupling is more difficult, primarily because shielding is much less effective for magnetic fields than for electric fields.²³ However, the principal method in dealing with both types of interference noise is proper physical layout and orientation of the conductors in the circuit.

Capacitive Coupling - The major source of capacitively coupled noise is digital circuits. Stray capacitances are usually very small and have large impedances at low frequencies. However, at the frequency of the fast edges of a digital signal, the impedance is much lower and significant feedthrough occurs, especially when the other conductor is a large impedance (Figure 41).




Figure 41. Capacitive Coupling

The principal method of dealing with this problem is to avoid placing digital paths close to critical signal paths. Special care must be taken to

avoid long stretches of parallel conductors and to avoid digital signals in the vicinity of high impedance nodes.

Shielding is a very effective method for eliminating capacitive coupling. Interposing a shield with zero impedance between the two conductors eliminates the problem. In this waveform generator, the signal conductors are confined to a small area around the four analog ICs. The physical separation of the conductors will be adequate for dealing with capacitive coupling and shielding will not be necessary - except possibly the final output. The only critical point is the hold capacitor of the deglitcher. Any high frequency noise injected on the output of the DAC either will be lost in the DAC's own transient as it switches, or will be long gone by the time the deglitcher acquires the output. The filter will remove any noise injected in its signal.

Magnetic Coupling - The situation that must be avoided here is loops with large areas in the vicinity of magnetic fields. The induced current in a loop is given by

$$V = 2\pi fAB\cos\beta \quad (48)$$

where f is the frequency of the flux density, A is the area of the loop, B is the field strength, and β is the angle between the field and the loop. The magnetic field strength can be reduced by increasing the distance between the loop and the source, but there is a limit to the reduction. The area of the loop should be kept small. This is done using twisted pair wiring and keeping conductors close together when the two form a loop. Also, long runs should be avoided. Susceptible signal paths should be run perpendicular to the field rather than parallel to minimize the coupling.²³

Shielding from magnetic fields is much less effective than shielding from electric fields because of limitations in materials available.²³ The primary method for reduction of magnetic coupling in this system is again proper attention to physical layout.

References

- ¹Ulrich L. Rhode, Digital PLL Frequency Synthesizers, (Englewood Cliffs, NJ: Prentice-Hall, Inc., 1983).
- ²Nasir T. Ahmed and T. Natarajan, Discrete Time Signals and Systems, (Reston, VA: Reston Publishing Company, Inc., 1983).
- ³Bruce A. Carlson, Communication Systems, (New York: McGraw-Hill, Inc., 1975).
- ⁴Benjamin C. Kuo, Automatic Control Systems, (Englewood Cliffs, NJ: Prentice-Hall, Inc., 1982).
- ⁵W. A. Evans and M. S. Towers, "Hybrid Technique in Waveform Generation and Synthesis", IEEE Proceedings, part G, 127(3):119-28.
- ⁶Allen Gersho, "Principles of Quantization", IEEE Transactions on Circuits and Systems, CAS-25(7):427-36, 1978.
- ⁷Daniele D. Caviglia et al, "Design and Construction of an Arbitrary Waveform Generator", IEEE Transactions on Instrumentation and Measurement, IM-32(3):398-403, 1983.
- ⁸Modules and Subsystems, vol II of 1984 Data Book, (Norwood, Mass.: Analog Devices).
- ⁹Eugene L. Zuch, ed, Data Acquisition and Conversion Handbook, (Mansfield, Mass.: Datel-Intersil, Inc., 1979).
- ¹⁰Richard C. Jaeger, "Tutorial: Analog Data Acquisition Technology Part I - Digital-to-Analog Conversion", IEEE Micro, p 20-37, May 1982.
- ¹¹David C. Pinkowitz, "A Designer's Guide to Deglitching DACs", Application Note, (Bohemia, NY: ILC Data Device Corporation).
- ¹²Alfi Moscovici, "Meet the Deglitcher", Application Note, (Bohemia, NY: ILC Data Device Corporation).
- ¹³Hal Chamberlin, Musical Applications of Microprocessors, (Rochelle Park, NJ: Hayden Book Company, Inc., 1980).
- ¹⁴Randal Thomas and Michael Jenkins, ed, Analog Switches and Their Applications, (Santa Clara, CA: Siliconix Inc., 1980).
- ¹⁵Arthur B. Williams, "Design of Active Filters Using Operational Amplifiers", Designer's Handbook of Integrated Circuits, ed Arthur B. Williams (New York: McGraw-Hill, Inc., 1984).
- ¹⁶John L. Hilburn and David E. Johnson, Manual of Active Filter Design, (2nd ed; New York: McGraw-Hill, Inc., 1983).

¹⁷Al Ryan and Tim Scranton, "D-C Amplifier Noise Revisited", Analog Dialogue, 18(7):3-10, 1984, (Bohemia, NY: Analog Devices, Inc.).

¹⁸George Erdi, "Amplifier Techniques for Combining Low Noise, Precision, and High-Speed Performance", IEEE Journal of Solid-State Circuits, SC-16(6):653-661, 1981.

¹⁹Donn Soderquist, "Minimization of Noise in Operational Amplifier Applications", Application Note, (Santa Clara, CA: Precision Monolithics Incorporated).

²⁰Linear and Conversion Products, 1984 Data Book, (Santa Clara, CA: Precision Monolithics Incorporated).

²¹Fanci F. Horton and Donald M. Hummels, "The H² Function Generator", (Design report, Kansas State University, Department of Electrical and Computer Engineering, 1983).

²²Raymond S. Turgel and N. Micheal Oldham, "High-Precision Audio-Frequency Phase Calibration Standard", IEEE Transactions on Instrumentation and Measurement, IM-27(4):460-4, 1978.

²³Henry W. Ott, Noise Reduction Techniques in Electronic Systems, (New York: John Wiley and Sons, 1976).

APPENDIX A

Circuit Diagram and Layout Details

Notes on the circuit diagram (Figure A1):

CMOS EPROMs - The chips are selected always ($CE = 0$) because the transition between read mode ($CE = 0$) and standby mode ($CE = 1$) creates large current spikes on the supply lines. The OE signals are used to select the chips. The time from OE high to output float is 130 ns while the time from OE low to output is 150 ns. Thus the output of one will go high-Z before the output of the other goes active.

Latches - CMOS provides symmetrical rise and fall times and HCMOS provides high speed to minimize data skew to the DAC.

Inverters - HCMOS is used to minimize any jitter in the latch controls.

Delay for DAC latches relative to the deglitcher control: 2 x inverter delay of 8 ns + latch delay of 16 ns + internal delay of DAC switches approximately 30 ns (assuming normal CMOS delay time) gives a total of 60 ns. The deglitcher delay is 30 ns so it will reach hold mode before the DAC output starts to change.

13 address bits are required for 8 kwords. This requires 4 MSI 4-bit adders. Bottom 13 bits are used - 12 for the ROM address and the 13th for the chip selects.

The ROMS are divided into two banks that are alternately selected - one for the first half of the signal (quadrants I and II) and one for the second half of the signal (quadrants III and IV).

Although not required, the top three bits of the adder output are connected to the latch inputs. The latch inputs must be connected somewhere so it might as well be the adder output to allow for future expansion. The

main reason the top three bits were not returned to the adder inputs was to avoid confusion in the circuit diagram.

18 bit words are required so 3 8-bit wide ROMS in parallel are used. The top 18 bits are sent to the DAC.

The DIP switch for manually setting the increment value is suitable initially but may be replaced with a latch for automated testing. The frequency selected is $I \times 6.25$ Hz, where I is the increment value.

A 50 kHz crystal is required. Using a higher frequency crystal and dividing it down is not a good idea due to noise considerations.

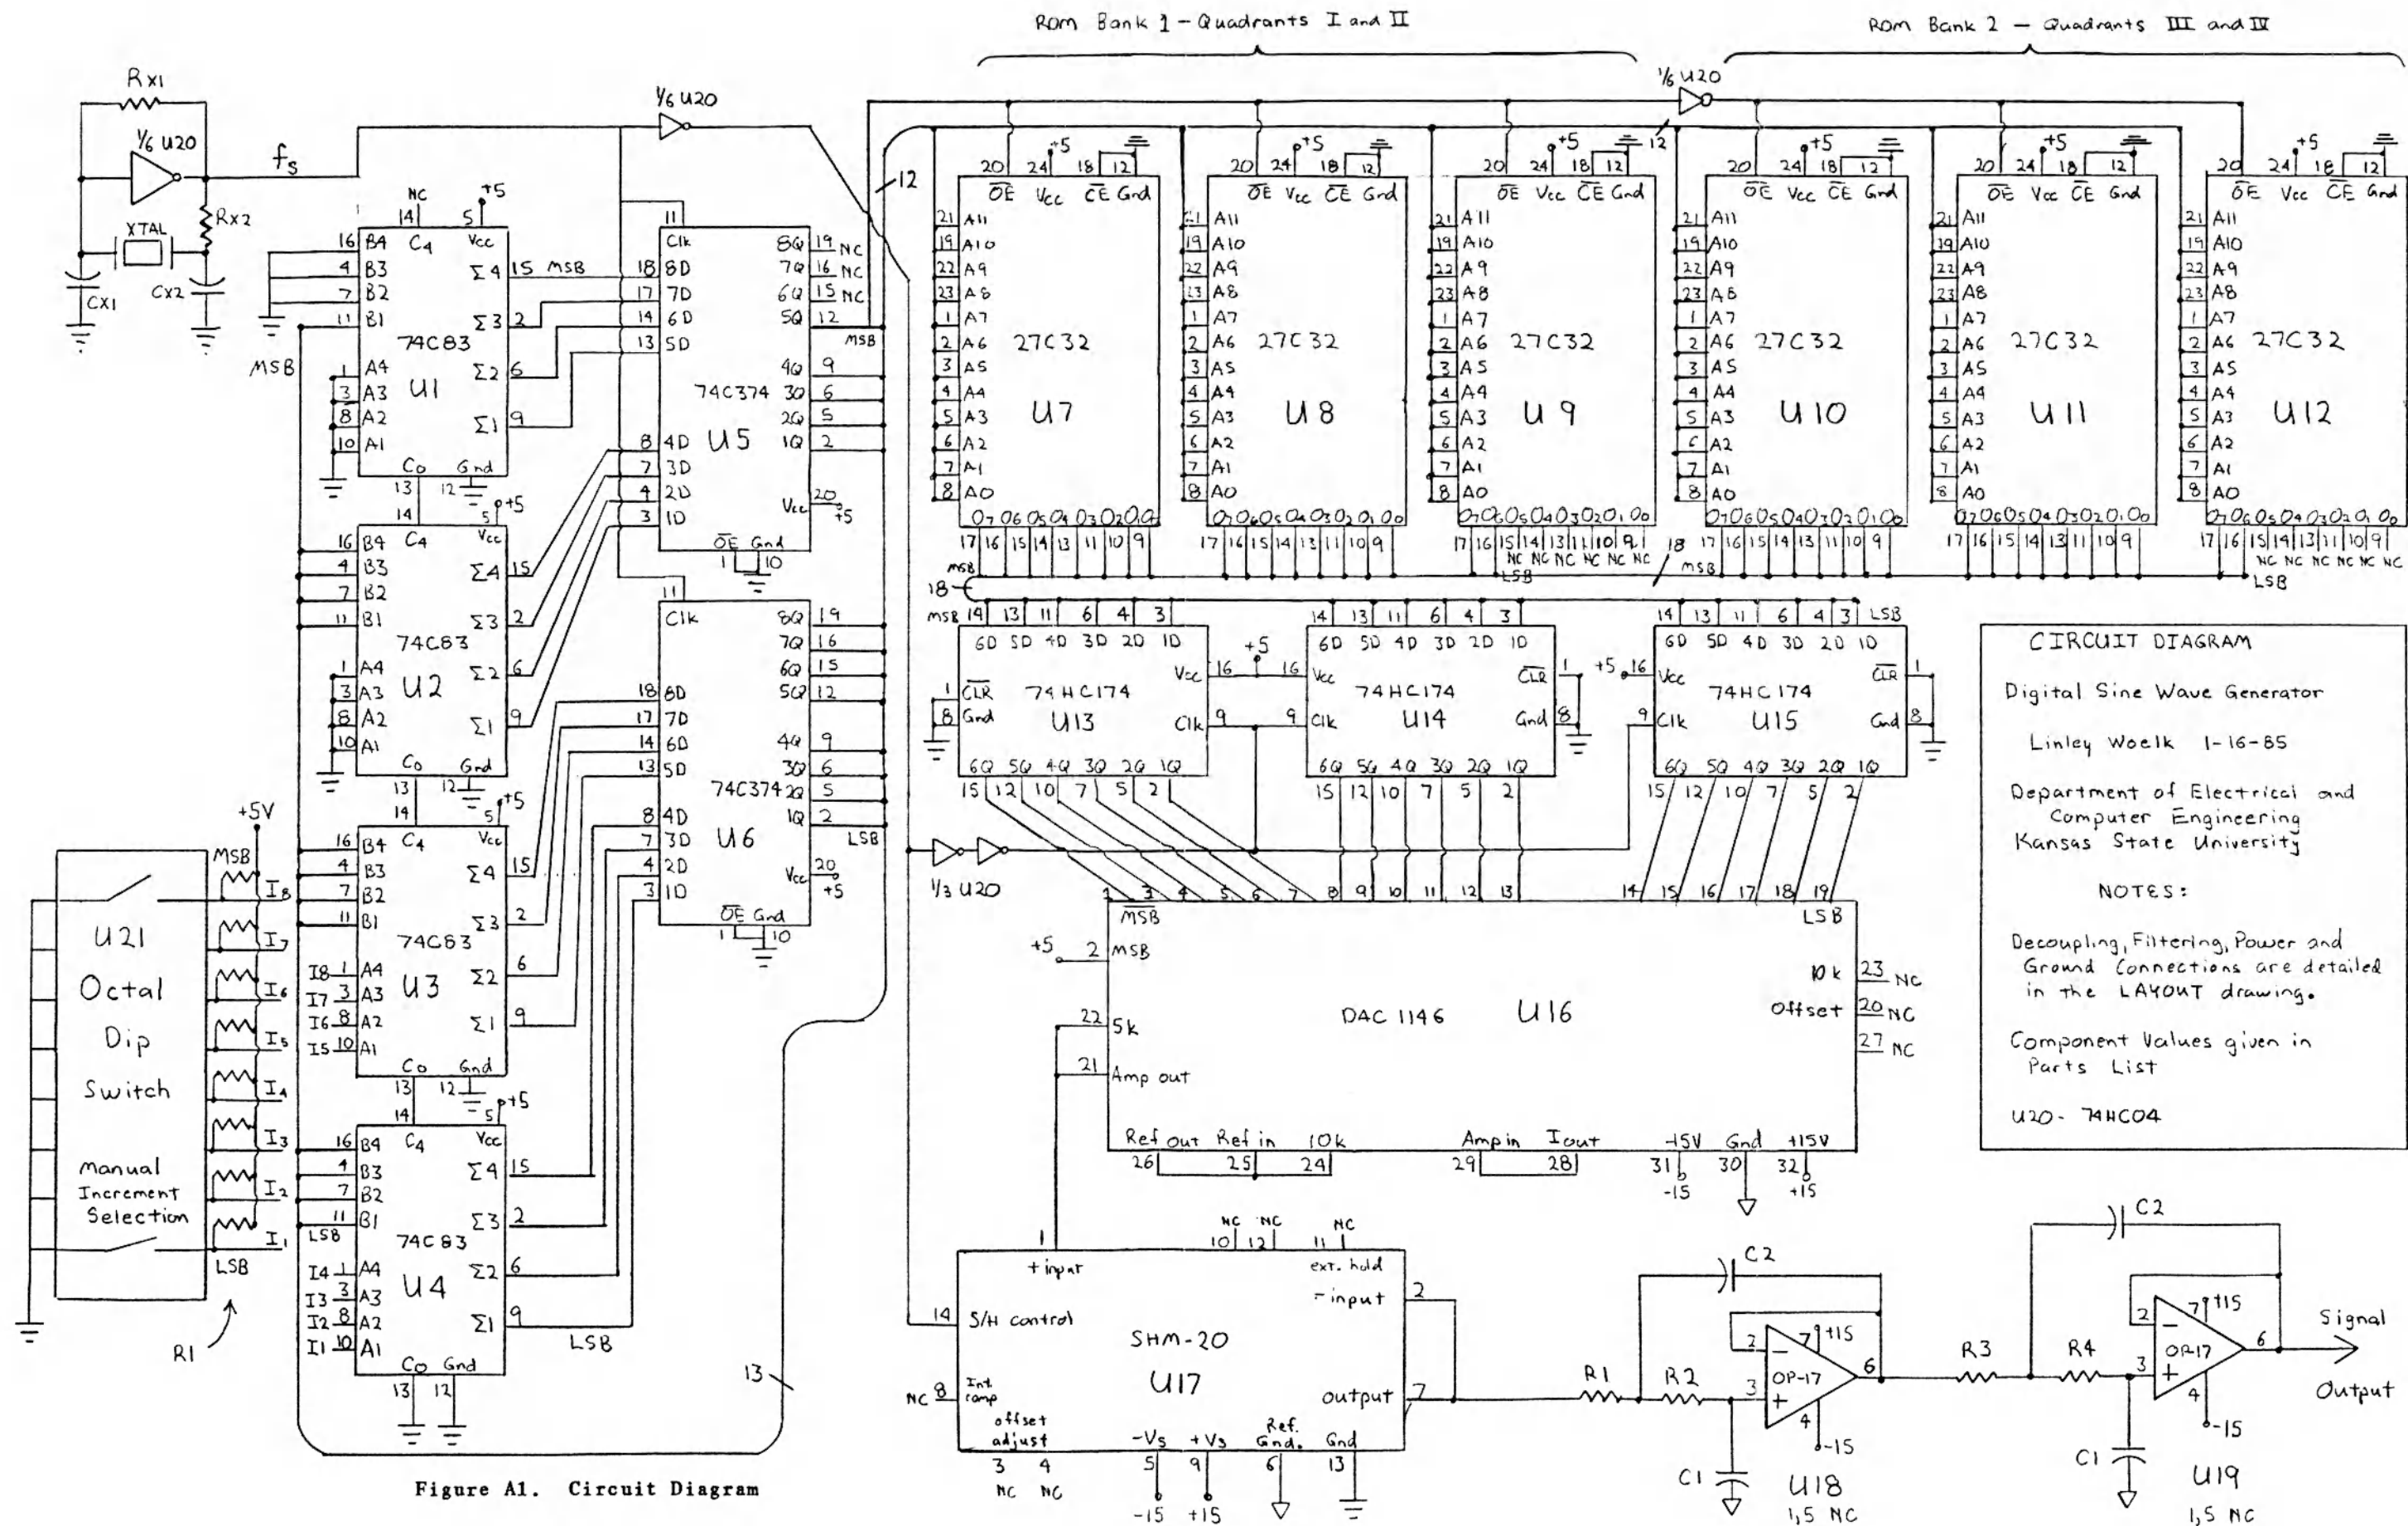


Figure A1. Circuit Diagram

Parts List

Number used	Designation	Part number and description
4	U1-U4	74C83 4-bit binary adders
2	U5,U6	74C374 octal D-FFs
6	U7-U12	27C32 CMOS EPROMS
3	U13-U15	74HC174 Hex D-FFs
1	U16	Analog Devices DAC 1146 18-bit DAC
1	U17	Datel Intersil SHM-20 Sample-Hold
2	U18,U19	PMI OP-17 precision JFET i/p op amp
2	C1	0.01 μ F polystyrene capacitor
2	C2	0.1 μ F polystyrene capacitor
1	R1	2.66 k Ω metal film resistors
1	R2	9.52 k Ω as close as possible to value
1	R3	887 Ω specified using series
1	R4	28.6 k Ω combinations of available values
10	R5	10 Ω metal film resistor
4	C3	100 μ F Tantalum capacitor
12	C4	0.1 ceramic capacitor




Figure A2. Layout diagram

- Notes: The single point ground is taken off board from only ONE connection
 → means return to single point by shortest possible route
 Keep all conductors as short as possible - analog conductors shown are critical.
 Decoupling capacitors should be as close as possible to the pins.
 NO analog or digital conductors (except S/H control) should cross the line of separation.

DIGITAL GENERATION OF LOW-FREQUENCY, LOW-DISTORTION TEST WAVEFORMS

by

LINLEY ELTON WOELK

B. S., Kansas State University, 1982

AN ABSTRACT OF A MASTER'S THESIS

submitted in partial fulfillment of the

requirements for the degree

MASTER OF SCIENCE

Department of Electrical and Computer Engineering

KANSAS STATE UNIVERSITY

Manhattan, Kansas

1985

Abstract

This paper describes the generation of high-purity waveforms by direct digital synthesis. The principle motivation for such a project comes from the requirements for dynamic testing of high-resolution analog-to-digital converters. Methods involving analyzing the frequency spectrum of the ADC's sampled data require a low-distortion test waveform so that the characteristics of the ADC are not obscured.

The design of a sine wave generator with equivalent accuracy to 16 bits is described. The underlying theory of sampling and waveform reconstruction is discussed and extended to include the various error sources encountered in real components. In particular, it is shown that nonlinearity errors must be included in the quantization error for accurate prediction of the expected signal-to-noise ratio of a digital waveform generator. The use of deglitchers and their related error sources are examined. It is shown that many of the errors normally of concern in sample-holds do not affect the signal integrity under certain conditions. It is also shown that an exponential rise is desirable for the staircase waveform while the normal high-speed settling characteristic of an amplifier is detrimental. Finally, some topics related to noise control in the system are discussed.